

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
26 April 2001 (26.04.2001)

PCT

(10) International Publication Number
WO 01/30035 A2

(51) International Patent Classification⁷: **H04L 12/56**

623 Kushland Way, Santa Cruz, CA 95064 (US). **BEYER, David, A.** [US/US]; 468 Paco Drive, Los Altos, CA 94024 (US).

(21) International Application Number: **PCT/US00/41180**

(22) International Filing Date: 16 October 2000 (16.10.2000)

(74) Agents: **FAHMI, Tarek, N. et al.**; Blakely, Sokoloff, Taylor & Zafman LLP, 12400 Wilshire Boulevard, 7th floor, Los Angeles, CA 90025 (US).

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
09/418,700 15 October 1999 (15.10.1999) **US**

(71) Applicant (for all designated States except US): **NOKIA WIRELESS ROUTERS, INC.** [US/US]; 313 Fairchild Drive, Mountain View, CA 94043 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

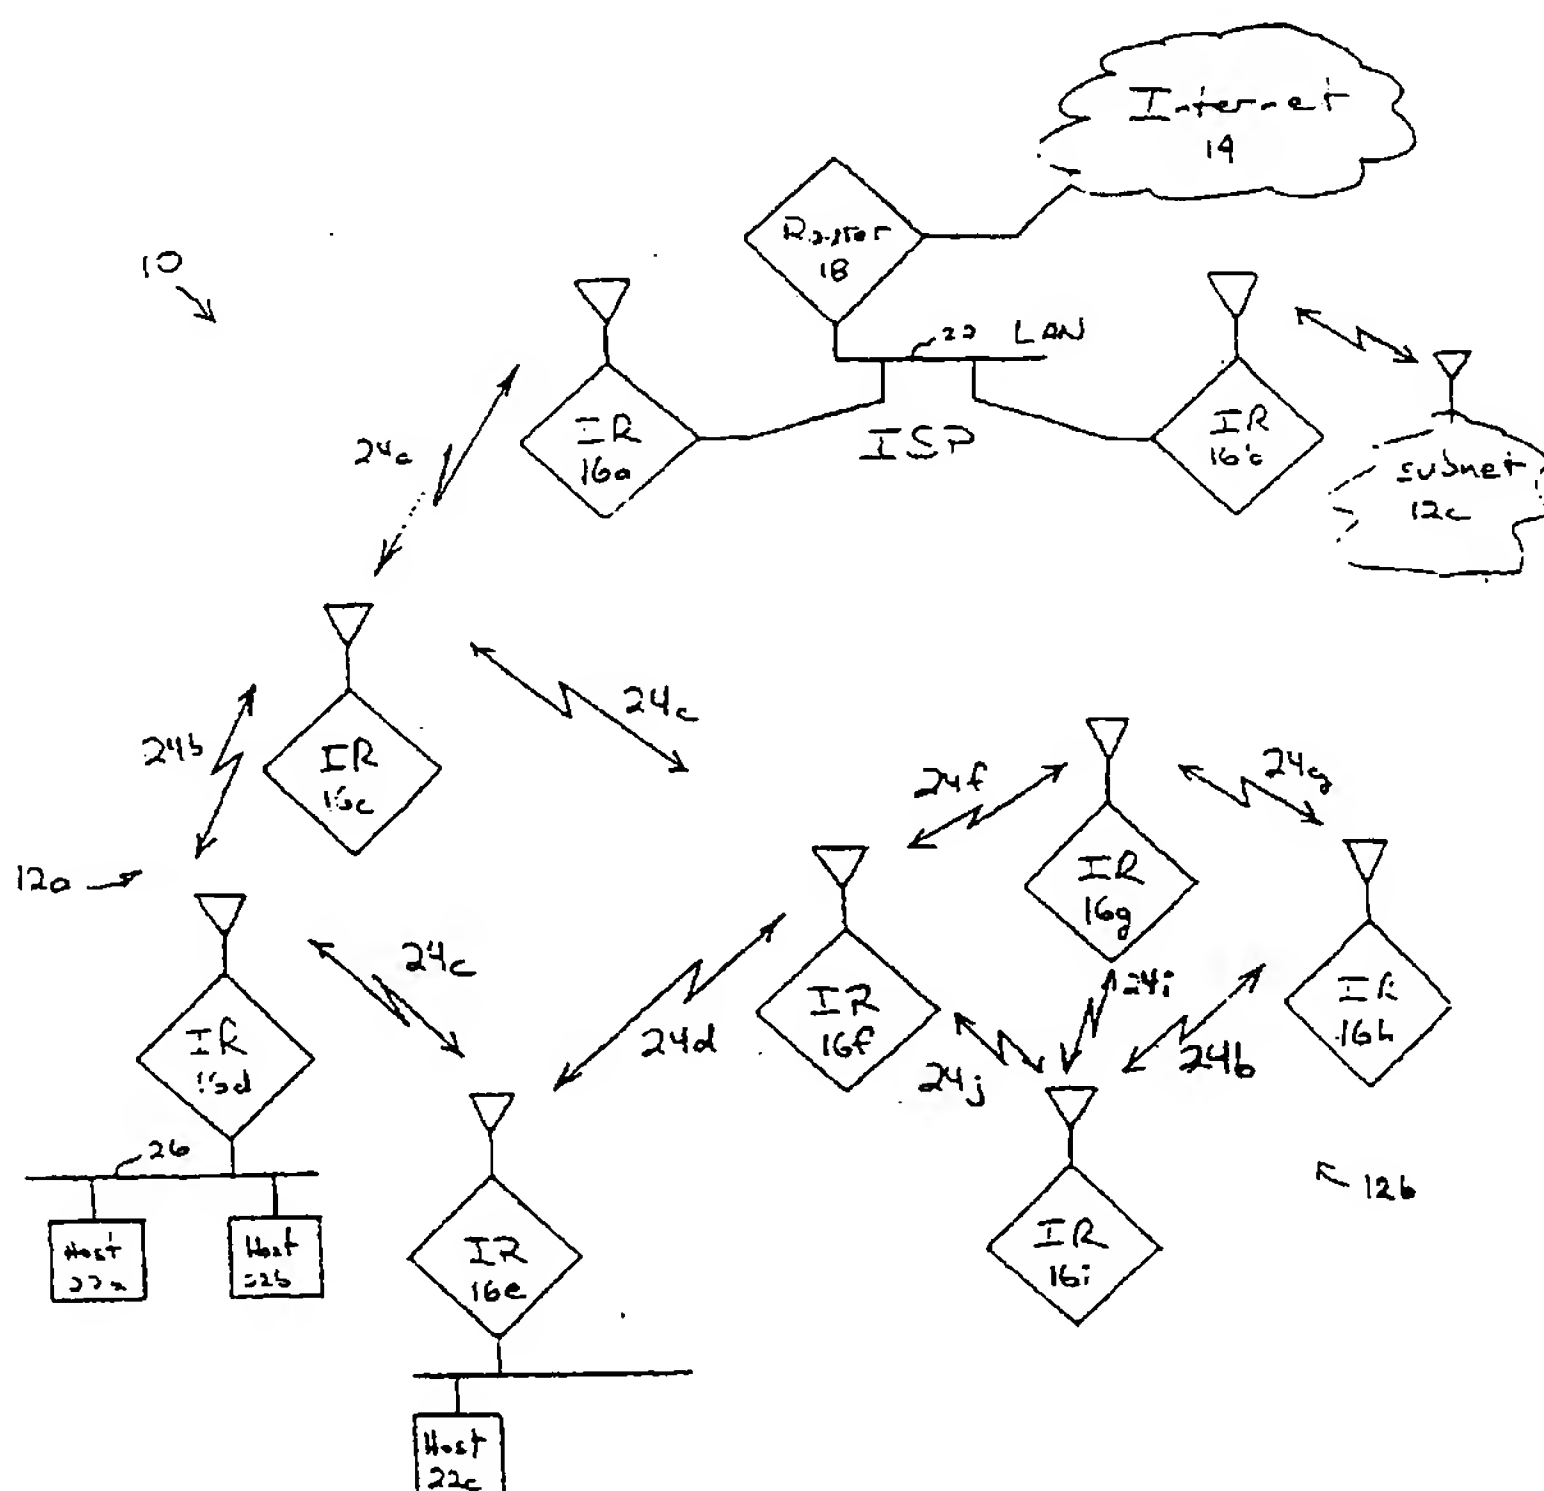
(72) Inventors; and

(75) Inventors/Applicants (for US only): **GARCIA-LUNA-ACEVES, J. Joaquin** [MX/US]; 82 Lakewood Circle, San Mateo, CA 94402 (US). **SPOHN, Marcelo** [BR/US];

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: **SYSTEM FOR COMMUNICATING LABELED ROUTING TREES TO ESTABLISH PREFERRED PATHS AND SOURCE ROUTES WITH LOCAL IDENTIFIERS IN WIRELESS COMPUTER NETWORKS**



(57) Abstract: One or more labeled routing trees (LRTs) are produced at a router of a computer network according to a shortest path determination made over a partial topology graph of the network, which graph is produced according to knowledge of adjacent links of the router and or more LRTs of neighboring routers. The LRTs of the router may be updated in response to receipt of routing state update messages, and such messages may include local link identifiers assigned by a head of a link to which the identifiers pertain, and node parameters of a tail of the link to which the local link identifiers pertain. The routing state update messages may be transmitted within the network: (i) in response to a new destination node being detected by an existing node within the network, (ii) in response to a destination becoming unreachable by a collection of the existing nodes, (iii) in response to the change in the cost of a path to at least one destination exceeding a threshold and/or (iv) in situations where a routing loop may be encountered among two or more of the nodes of the network (e.g., at times when a path implied in the LRT of the router leads to a loop).



Published:

- *Without international search report and to be republished upon receipt of that report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

**SYSTEM FOR COMMUNICATING LABELED ROUTING TREES TO ESTABLISH
PREFERRED PATHS AND SOURCE ROUTES WITH LOCAL IDENTIFIERS IN WIRELESS
COMPUTER NETWORKS**

STATEMENT OF GOVERNMENT LICENSE RIGHTS

The United States Government has a paid-up license in portions of this invention and the right in limited circumstances to require the patent owner to license others on reasonable terms as provided for by the terms of Contract No.: DAAH01-98-C-R005, awarded by the U.S. Army Aviation & Missile Command.

FIELD OF THE INVENTION

The present invention relates to routing protocols in computer networks, and more particularly, routing protocols in ad hoc networks with radio links.

BACKGROUND

Multi-hop packet-radio networks, or ad hoc networks, consist of mobile hosts interconnected by routers that can also move. The deployment of such routers is ad hoc and the topology of such networks is very dynamic, because of host and router mobility, signal loss and interference, and power outages. In addition, the channel bandwidth available in ad hoc networks is relatively limited compared to wired networks, and untethered routers may need to operate with battery-life constraints. In these networks, routing must preferably be accomplished using as few a number of control messages and neighbor-to-neighbor handshakes as possible, in order to conserve channel bandwidth for user data and preserve the battery life of untethered nodes. Because of the dynamics of the topology in an ad hoc network, broadcast radio links are preferable for interconnecting routers without the need for topology planning.

Routing protocols for computer networks can be categorized according to: (a) the type of information the protocols use to compute preferred paths, and (b) the way in which routers obtain routing information. In terms of the type of information used by routing protocols, routing protocols can be classified into link-state protocols and

distance-vector protocols. Routers running a link-state protocol use topology information to make routing decisions; routers running a distance-vector protocol use distances and, in some cases, path information, to destinations to make routing decisions.

In terms of the way in which routers obtain information, routing protocols have been classified as either table-driven or on-demand. In an on-demand routing protocol, routers maintain path information for only those destinations that they need to contact as a source or relay of information. The basic approach consists of allowing a router that does not know how to reach a destination to send a flood-search message to obtain the path information it needs. One of the first routing protocols of this type was proposed to establish virtual circuits in the MSE network, see V.O.K. Li and R. Chang, "Proposed routing algorithms for the US Army mobile subscriber equipment (MSE) network," Proc. IEEE MILCOM'86, Monterey, California, October 1986, and there are several other, recent examples of this approach (e.g., AODV, see C. Perkins, "Ad Hoc On Demand Distance Vector (AODV) Routing," draft-ietf-manet-aodv-00.txt, November 1997; ABR, see C-K. Toh, "Wireless ATM & Ad Hoc Networks," Kluwer, November 1996; DSR, see D. Johnson and D. Maltz, "Protocols for Adaptive Wireless and Mobile Networking," IEEE Pers. Commun., Vol. 3, No. 1, February 1996; TORA, see V. Park and M. Corson, "A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks," Proc. IEEE INFOCOM 97, Kobe, Japan, April 1997; SSA, see R. Dube et al., "Signal Stability-Based Adaptive Routing (SSA) for Ad Hoc Mobile Networks," IEEE Pers. Commun., February 1997; and ZRP, see Z. Haas and M. Pearlman, "The Zone Routing Protocol for Highly Reconfigurable Ad Hoc Networks," Proc. ACM SIGCOMM 98, Vancouver, British Columbia, August 1998). The Dynamic Source Routing (DSR) protocol has been shown to outperform many other on-demand routing protocols. J. Broch et al, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," Proc. ACM MOBICOM 98, Dallas, TX, October 1998. The existing on-demand routing protocols differ on the specific mechanisms used

to disseminate flood-search packets and the responses thereto, the means used to cache information received during other nodes' searches, and the manner in which to determine the cost of a link and the existence of a neighbor. However, one common characteristic of all of the on-demand routing protocols reported to date is that such protocols are based on distances to destinations. Stated differently, there have been no on-demand link-state routing protocol proposals to date.

In a table-driven scheme, each router maintains path information for each known destination in the network and updates its routing-table entries as needed. Examples of table-driven algorithms based on distance vectors are the routing protocol of the DARPA packet-radio network, J. Jubin and J. Tornow, "The DARPA Packet Radio Network Protocols," Proceedings of the IEEE, Vol. 75, No. 1, January 1987; DSDV, C. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers," Proc. ACM SIGCOMM 94, London, UK, October 1994; WRP, S. Murthy and J.J. Garcia-Luna-Aceves, "An Efficient Routing Protocol for Wireless Networks," ACM Mobile Networks and Applications Journal, Special issue on Routing in Mobile Communication Networks, 1996; WIRP, J.J. Garcia-Luna-Aceves et al., "Wireless Internet Gateways (WINGS)," Proc. IEEE MILCOM'97, Monterey, California, November 1997; and least-resistance routing protocols, M Pursley and H.B. Russell, "Routing in Frequency-Hop Packet Radio Networks with Partial-Band Jamming," IEEE Trans. Commun., Vol. 41, No. 7, pp. 1117-1124, 1993.

Prior table-driven approaches to link-state routing in packet-radio networks are based on topology broadcasts. However, disseminating complete link-state information to all routers incurs excessive communication overhead in an ad hoc network because of the dynamics of the network and the small bandwidth available. Accordingly, all existing link-state routing approaches for packet-radio networks have been based on hierarchical routing schemes. R. Ramanathan and M. Steenstrup, "Hierarchically-organized, Multihop Mobile Wireless Networks for Quality-of-Service Support," ACM Mobile Networks and Applications, Vol.3, No. 1, pp. 101-119, 1998; C.V.

Ramamoorthy and W. Tsai, "An Adaptive Hierarchical Routing Algorithm," Proceedings of IEEE COMPSAC '83, Chicago, Illinois, pp. 93-104, November 1983; and M. Steenstrup (ed.), Routing in Communication Networks, Prentice-Hall, 1995. Also, prior proposals for link-state routing using partial link-state data without clusters, see, e.g., J.J. Garcia-Luna-Aceves and J. Behrens, "Distributed, scalable routing based on vectors of link states," IEEE Journal on Selected Areas in Communications, Vol. 13, No. 8, 1995; and J.J. Garcia-Luna-Aceves and M. Spohn, "Scalable Link-State Internet Routing," Proc. IEEE International Conference on Network Protocols (ICNP 98), Austin, Texas, October 14-16, 1998, require routers to explicitly inform their neighbors which links they use and which links they stop using.

A number of prior routing protocols are based on the notion of routing trees, in which routers communicate either the state (i.e., cost or length) of the links in a shortest-path routing tree, or the distance from the root of the tree and the second-to-last hop in the routing tree for each node in the tree. An early example of this type of protocol was proposed in U.S. Patent 4,466,060 to Riddle. In Riddle's protocol, a router communicates different routing trees to different neighbors; such trees are called "exclusionary trees" and specify the preferred paths to destinations excluding those paths that involve the router to which the update is being sent. An update packet or message specifies an entire exclusionary tree. Another protocol based on routing trees was reported by Garcia-Luna-Aceves, J.J. Garcia-Luna-Aceves, "A Fail-Safe Routing Algorithm for Multihop Packet-Radio networks," Proc. IEEE Infocom 86, Miami, FL, April 1986, which protocol differs from Riddle's protocol in that the same routing tree is sent incrementally by a router to all its neighbors. A. Humblet, "Another Adaptive Shortest-Path Algorithm," IEEE Trans. Comm., Vol.39, No.6, June 1991, pp.995--1003 (and see US Patent 4,987,536); Cheng et al., C. Cheng et al., "A Loop-Free Extended Bellman-Ford Routing Protocol without Bouncing Effect", Proc. ACM SIGCOMM 89, pp.224--236; B. Rajagopalan and M. Faiman, "A Responsive Distributed Shortest-Path Routing Algorithm within Autonomous Systems," Journal of Internetworking: Research and Experience}, Vol. 2, No. 1, March 1991, pp. 51-69; and S. Murthy and J.J. Garcia-

Luna-Aceves, "Loop-Free Internet Routing Using Hierarchical Routing Trees," Proc. IEEE INFOCOM 97, Kobe, Japan, April 7-11, 1997, have all proposed protocols based on source trees in which a router communicates to its neighbors the same shortest-path routing tree incrementally and these latter examples differ from the above-cited protocol by Garcia-Luna-Aceves in the way in which a router obtains its own source tree from the trees reported by its neighbors. An important feature of all prior routing tree-based protocols is the fact that they are all based on communicating routing trees in terms of the lengths of the links and the identifiers of the nodes that form part of the routing trees.

SUMMARY OF THE INVENTION

In one embodiment of the present scheme, one or more labeled routing trees (LRTs) are produced at a router of a computer network according to a shortest path determination made over a partial topology graph of the network, which graph is produced according to knowledge of adjacent links of the router and one or more LRTs of neighboring routers. The LRTs of the router may be updated in response to receipt of routing state update messages, and such messages may include local link identifiers assigned by a head of a link to which the identifiers pertain, and node parameters of a tail of the link to which the local link identifiers pertain. The routing state update messages may be transmitted within the network: (i) in response to a new destination node being detected by an existing node within the network, (ii) in response to a destination becoming unreachable by a collection of the existing nodes, (iii) in response to the change in the cost of a path to at least one destination exceeding a threshold delta, and/or (iv) in situations where a routing loop may be encountered among two or more of the nodes of the network (e.g., at times when a path implied in the LRT of the router leads to a loop).

In another embodiment, the present routing protocol allows for distributing local link identifiers among nodes of a computer network within routing state update messages. These local link identifiers are preferably assigned by a head of a link to

which the identifiers pertain. The routing state update messages may include state parameters for nodes at tails of links to which the identifiers pertain. The routing protocol may further provide for distributing labeled routing trees of the nodes of the computer network among the nodes, for example wherein each node of the computers network may distribute its labeled routing trees to its neighboring nodes.

In such a routing protocol each node of the computer network preferably maintains one labeled routing tree per type of service offered in the network. These labeled routing trees may be generated according to a partial topology graph derived from the node's adjacent links and labeled routing trees of neighboring nodes, for example by applying a path selection algorithm to the partial topology graph. The labeled routing trees may then be updated in response to receipt of the routing state update messages. Such updates may be made according to either an optimum routing approach or a least overhead routing approach, depending upon which approach is used within the computer network.

For the optimum routing approach, routing state update messages are preferably transmitted: (i) when a routing state update message which causes a link to be added to a subject node's labeled routing trees is received, (ii) when a more recent routing state update message regarding an existing link in the subject node's labeled routing trees is received, and/or (iii) when a destination becomes unreachable by the subject node. In the least overhead routing approach, routing state update messages may be transmitted: (i) when a subject node finds a new destination or any neighbors of the subject node report same, (ii) when a destination becomes unreachable by the subject node or any of its neighbors, (iii) when the subject node finds that the change in the cost of a path to at least one destination exceeds a threshold delta, (iv) when a path implied in any of the labeled routing trees of the subject node leads to a loop, and/or a new successor chosen to a given destination has an address larger than an address of the subject node and a reported distance from the new successor to a particular destination is larger than a reported distance from a previous successor to that particular destination. Another embodiment provides a routing update message that includes information regarding

performance characteristics and addressing information for a link of a computer network and a node of the network at a tail end of the link. In some cases, the routing update may further include a time stamp assigned by a node at a head end of the link, a type of service vector, and addressing information for the node at a head end of the link. The performance characteristics preferably include link state parameters (e.g., link delay, link cost, available bandwidth, and/or reliability) of the link, while the type of service vector preferably specifies a type of service routing tree in which the link is being used by a node transmitting the routing update message. The addressing information of the link is preferably specified in the form of a local link identifier assigned to the link by the node at a head end of the link.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements and in which:

Figure 1 illustrates an example of an ad hoc wireless network with routers configured in accordance with the present routing protocol;

Figures 2a - 2d illustrate an example of a six-node wireless network with routers configured in accordance with the present routing protocol and the labeled routing trees developed at various ones of the routers;

Figures 3a - 3f illustrate an example of a six-node wireless network with routers configured in accordance with the present routing protocol and the routing state updates generated after certain link failures within the network; and

Figures 4a - 4f illustrate an example of how a link failure within a wireless network made up of a number of routers configured in accordance with the present routing protocol may not lead to the generation of new update messages by such routers when all such routers still have a path to all available destinations in the network.

DETAILED DESCRIPTION

A scheme for enabling routing of data packets in a computer network along preferred paths either on a hop-by-hop basis, or by specifying a source route efficiently by means of local identifiers is disclosed herein. Although discussed with reference to certain illustrated embodiments, upon review of this specification, those of ordinary skill in the art will recognize that the present scheme may find application in a variety of systems. Therefore, in the following description the illustrated embodiments should be regarded as exemplary only and should not be deemed to be limiting in scope.

In the present scheme, each packet being routed carries a routing operation code (ROC) instructing the receiving router which routing method to apply to forward the packet. A packet can be forwarded in any of the following forwarding modes: (a) a conventional hop-by-hop routing mode, which uses the routing table of the router; or (b) a source-routing mode, in which the entire source route is specified in the packet using local link identifiers instead of the addresses of relay routers as is common in prior source routing approaches. These forwarding modes are enabled using the present routing protocol, termed the Adaptive Internet Routing (AIR) protocol, which allows for the dissemination of link-state information and node-state information in the form of labeled routing trees (LRTs).

With AIR, a router sends updates to its neighbors regarding the links and nodes in its preferred paths to destinations. The links and nodes along the preferred paths from a source to each desired destination constitute an LRT that implicitly specifies the complete paths from the source to each destination. Each link is labeled with a local link identifier (LLID).

Each link in an LRT is directed and has a head-of-link node and a tail-of-link node. The head of the link labels the link with an LLID that differentiates that link from all other links from the same node to other neighbors. The LLID of a link is much smaller (in terms of the number of bits required to specify the LLID) than the address of a node.

Each router maintains an LRT for each type of service defined in the network (e.g., minimum-hop, minimum delay, and maximum bandwidth paths of the smallest number of hops). Each router also maintains a routing or topology graph that includes the state information about adjacent links and the LRTs reported by its neighbors. Each router computes its LRT for a given type of service based on the link and node information available in its topology graph.

A router reports changes to any of its LRTs to all its neighbors incrementally or atomically. The rules used to decide when a router should communicate changes to the state of a node or link can be based on optimum routing or least overhead routing approaches. The aggregation of adjacent links and routing trees reported by neighbors then constitutes the partial topology known by a router. To minimize the number of times the state of a link is communicated to neighbors, a router may assign a type vector to each link, with each bit of the vector indicating in which LRT the link is being used by the router.

AIR does not require backbones, the dissemination of complete cluster topology within a cluster, or the dissemination of the complete inter-cluster connectivity among clusters. Furthermore, AIR can be used with distributed hierarchical routing schemes proposed in the past for either distance-vector or link-state routing. See, e.g., L. Kleinrock and F. Kamoun, "Hierarchical Routing for Large Networks: Performance Evaluation and Optimization," *Computer Networks*, Vol. 1, pp. 155-174, 1977; M. Steenstrup (ed.), *Routing in Communication Networks*, Prentice-Hall, 1995; S. Murthy and J.J. Garcia-Luna-Aceves, "Loop-Free Internet Routing Using Hierarchical Routing Trees," *Proc. IEEE INFOCOM 97*, Kobe, Japan, April 7-11, 1997; and J. Behrens and J. J. Garcia-Luna-Aceves, "Hierarchical Routing Using Link Vectors," *Proc. IEEE INFOCOM 98*, San Francisco, California, March 29-April 2, 1998.

A router chooses its preferred paths for a given type of service using a local path-selection algorithm. One preferred path-selection algorithm for use according to the present scheme is a modification of the shortest-path first (SPF) algorithm. The result of running the path-selection algorithm over the topology graph is an LRT for a

given type of service that specifies, for each node in the LRT: the address of the head of the link incident to the node, the state parameters of the link, the state parameters of the node, and the LLID assigned to the link by the head of the link. Of course, other path selection algorithms may be used to provide similar outputs.

From its LRT, a router can compute a source route to a given destination. Because links are labeled with LLIDs assigned by the head of the links, a router can uniquely specify a source route to a destination using its own address followed by a sequence of LLIDs corresponding to its preferred path to the destination, rather than as a sequence of much larger network or link-level addresses. Although source routing and the use of local link identifiers have been used independently of one another in routing and bridging protocols in the past, AIR is the first routing protocol that disseminates the LLIDs of links, thus allowing a compact specification of source routes, making much more efficient use of the available bandwidth.

The present routing scheme is thus well suited for an ad hoc network that provides a seamless extension of the Internet Protocol (IP) to the ad hoc wireless environment. AIR will be described in terms of its operation in internet radios or IRs, which are wireless routers. However, it will be evident to those of ordinary skill in the art that AIR applies to computer networks and inter-networks that need not be based on wireless links for router interconnection. Figure 1 illustrates aspects of an exemplary ad hoc internet that will assist in understanding the remaining discussion.

Ad hoc network 10 may be considered as a number of sub-networks 12a, 12b, 12c, which provide an extension of the Internet 14 through a number of Internet Radios or IRs 16a-16i. Each IR 16a-16i is a wireless router with an assigned IP address and a medium access control (MAC) address. In general, IRs 16a-16i operate over one or a multiplicity of radio channels using spread-spectrum wireless communication techniques common in the art. For example, the IRs 16a-16i may operate in one of the unregulated UHF frequency bands, thereby obviating the need for operating licenses. Coupling of ad hoc network 10 to the Internet 14 is achieved through a router 18, which may be operated by an Internet Service Provider (ISP). As shown, a single ISP may

operate a LAN 20 to which multiple IRs are connected. In such a scheme, IRs 16a and 16b may act as "AirHeads", providing gateway service to Internet 14 via router 18.

Some IRs, e.g., IRs 16d and 16e of Figure 1, may be associated with hosts, 22a, 22b and 22c, which can be accessed by any Internet user through ad hoc network 10. Like any router, each IR 16a-16i processes all messages, changes in the cost of an adjacent link, adjacent-link failures, and new-neighbor notifications one at a time and in the order in which it detects them.

Any IR 16a-16i in Figure 1 can consider another IR to be adjacent (we call such an IR a "neighbor") if there is radio connectivity between the two IRs and one IR, e.g., IR 16g, can receive and acknowledge packets from the other IR, e.g., IR 16h.

Accordingly, a physical broadcast link connecting multiple IRs is mapped into multiple point-to-point bidirectional links defined for the same IRs. Each pair of adjacent IRs defines two point-to-point bidirectional links between them, one in each direction. Each point-to-point bidirectional link has a head node of the link and a tail node of the link.

The present routing scheme can be brought to practice together with the methods described in co-pending Application No. 09/248,738, entitled "Adaptive Communication protocol for Wireless Networks," filed February 10, 1999, and assigned to the Assignee of the present invention, incorporated herein by reference, for the assignment of logical link identifiers by one node to each of its one-hop neighbors. However, it will be evident to those of ordinary skill in the art that the present scheme can make use of any link-level service that enables a router to use a local identifier for each of its one-hop neighbors. The description of the present routing scheme thus assumes the existence of local identifiers for the links between a router and its immediate neighbors.

An underlying protocol, the above-noted neighbor protocol, assures that each IR 16a-16i detects within a finite time the existence of a new neighbor IR and the loss of connectivity with a neighbor IR. The neighbor protocol assumed in the present scheme can be brought to practice using link-layer retransmission strategies common in the art.

1. AIR Operation

In AIR, each IR reports to its neighbors the characteristics of every link and IR it uses to reach a destination in the ad hoc network. The set of links and IRs used by an IR in its preferred path to destinations is called the LRT of the IR. If multiple types of service (TOS) are defined in the network, an IR maintains one LRT per TOS. An IR therefore knows its adjacent links and the LRTs reported by its neighbors for each TOS, and the aggregation of an IR's adjacent links and the LRTs reported by its neighbors constitutes a partial topology graph. The links in the LRT and topology graph should be adjacent links or links reported by at least one neighbor. An IR may use the topology graph to generate its own LRT for each TOS. Each IR derives its LRT and routing table specifying the successor to each destination for each TOS by running a local path-selection algorithm for each TOS on its topology graph.

An IR communicates the updates it makes to its routing tree for any TOS. Because each IR communicates its routing tree for each TOS to its neighbors, the deletion of a link no longer used to reach a destination for a given TOS is implicit with the addition of the new link used to reach the destination and need not be sent explicitly as an update; an IR makes explicit reference to a failed link only when the deletion of a link causes the IR to have no paths to one or more destinations, in which case the IR cannot provide new links to make the deletion of the failed link implicit.

The basic update unit used in AIR to communicate changes to routing trees is the routing-state update (RSU), which describes the state of a link and the node at the end of the link. The head node of a link is the only IR that can report changes in the parameters of that link. RSUs are validated using sequence numbers, and each IR erases a link from its topology graph if the link is not present in the routing trees of any of its neighbors.

The operation of AIR can be described formally by pseudocode. As with any such description, however, there are various equivalent formulations that, upon review of the code provided, will be evident as being equivalent thereto to someone of ordinary

skill in the art. The pseudocode description is set forth at the end of this detailed description, and the following variables and expressions are used therein:

T : a constant defining the number of Types of Service (TOS) in the network

$\Delta[T]$: $\Delta[0], \Delta[1], \dots, \Delta[T-1]$ is a vector corresponding to thresholds to changes in the cost of a path to a destination for a given TOS.

AGEOUT_INTERVAL: a constant defining the time a failed link stays in the Topology Graph before being aged-out

ORA: a constant, TRUE if AIR is running under the Optimum Routing Approach

LORA: a constant, TRUE if AIR is running under the Least Overhead Routing Approach

TG_i : Topology Graph at router i

LRT_i : Current Labeled Routing Tree at router i

LRT_i' : Labeled Routing Tree created the last time an update message was generated by i

N_i : set of neighbors of router i

NS_i : set to TRUE if i has not sent its Labeled Routing Tree to a neighbor

M_i : set to TRUE if i must report changes to the Labeled Routing Tree

T_i : system time used to generate time stamps to RSU_i

$(u, v, llid, t; l[T], \{n\}, tos[T], del)$: An entry for link (u, v) in TG_i, LRT_i, TG_k^i , and LRT_k^i , where $k \in N_i$ and $\{l\} = \{llid, l[T], tos[T], del\}$.

u : Network address of the head of the link

v : Network address of the tail of the link

$llid$: Local link identifier

t : timestamp

$l[T]$: $l[0], l[1], \dots, l[T-1]$ is a vector corresponding to the performance parameters of the link. The parameter $l[0]$ corresponds to the cost of the link, the remaining parameters may be delay, bandwidth, and reliability of the link, etc.

$\{n\}$: The state parameters of router v

tos[T]: tos[0], tos[1], . . . , tos[T - 1] is the Type of Service bit-vector. A bit x is set to 1 when the link is added to the Labeled Routing Tree with TOS x.

del: Set to TRUE if the link cannot be used in the computation of the Labeled Routing Trees

(d[T], pred[T], suc[T], d'[T], d''[T], suc'[T], nbr): Variables assigned to a vertex v in TG_i ,

LRT_i , TG_k^i , and LRT_k^i , where $k \in N_i$

d[T]: Cost of the path $i \rightarrow v$ in Labeled Routing Tree x; $\forall x \in [0; T - 1]$

1] pred[T]: Predecessor (link) of vertex v in Labeled Routing Tree x; $\forall x \in [0; T - 1]$

suc[T]: Next-hop towards vertex v in Labeled Routing Tree x; $\forall x \in [0; T - 1]$

d'[T]: Previous distance to v reported by suc'[T]

d''[T]: Cost of the path $i \rightarrow v$ the last time the cost of the path changed by $\Delta[T]$

1] suc'[T]: Previous successor towards v in Labeled Routing Tree x; $\forall x \in [0; T - 0]$

be nbr: Set to i if the cost of the path to v has increased but no update message must be generated

(u, v, llid, t, l[T], {n}, tos[T]): Entry in an update message (RSU)

u: Network address of the head of the link

v: Network address of the tail of the link

llid: Local link identifier

t: Time stamp assigned to RSU

l[T]: Vector corresponding to the performance parameters of the link

{n}: The state parameters of router v

tos[T]: Type of Service bit-vector

The pseudocode specifies the main procedures of AIR used to update the routing table and the link-state database at a router i for both an Optimum Routing Approach

(ORA) and a Least Overhead Routing Approach (LORA). Procedure NodeUp is executed when a router i starts up. The neighbor set of the router is empty initially.

If the neighbor protocol reports a new link to a neighbor k (procedure {lem NeighborUp}), the router then runs Update with the appropriate message as input; the RSU in the message gets a current time stamp. The same approach is used for link failures (NeighborDown) and changes in the parameters of a link (LinkChange). When a router establishes connectivity to a new neighbor, the router sends its complete labeled routing tree to the neighbor (much like a distance vector protocol sends its complete routing table). The RSUs that are to be broadcast to all neighbors are inserted into MSG_i .

The procedure Update is executed when router i receives an update message from neighbor k or when the parameters of an outgoing link have changed. First, the topology graphs TG_i and TG_k^i are updated, then the labeled routing trees LRT_k^i and LRT_i are updated, which may cause the router to update its routing table and to send its own update message.

The state of a link in the topology graph TG_i is updated with the new parameters for the link if the routing-state update in the received message is valid, i.e., if the RSU has a larger time stamp than the time stamp of the link stored in TG_i .

The parameters of a link in TG_k^i are always updated when processing an RSU sent by a neighbor k , even if the link-state information is outdated, because they report changes to the labeled routing tree of the neighbor. A node in a labeled routing tree LRT_k^i for a given TOS can have only one link incident to it. Hence, when i receives an RSU for link (u,v) from k the current incident link (u', v) to v , $u \neq u'$, is deleted from TG_k^i .

The information of an RSU reporting the failure of a link is discarded if the link is not in the topology graph of the router.

A shortest-path algorithm (SPF) based on Dijkstra's SPF (procedure BuildShortestPathTree) is run on the updated topology graph TG_k^i to construct a new

labeled routing tree LRT_k^i , and then run on the topology graph TG_i to construct a new labeled routing tree LRT_i .

The incident link to a node v in router's i new labeled routing tree is different from the link in the current labeled routing tree LRT_i only if the cost of the path to v has decreased or if the incident link in LRT_i was deleted from the labeled routing trees of all neighbors.

A new labeled routing tree $newLRT$ for a neighbor k , including the router's new labeled routing tree, is then compared to the current labeled routing tree LRT_k^i (procedure `UpdateNeighborTree`), and the links that are in LRT_k^i but not in $newLRT$ are deleted from TG_k^i . After deleting a link (u, v) from TG_k^i the router sets $TG_i(u, v).del$ to `TRUE` if the link is not present in the topology graphs TG_x^i , $\forall x \in N_i$.

If a destination v becomes unreachable, i.e., there is no path to v in the new labeled routing tree $newLRT$, then RSUs will be broadcast to the neighbors for each link in the topology graph TG_i that have v as the tail node of the link and a link cost infinity.

This specification assumes that the Link Layer provides reliable broadcast of network-level packets and consequently update messages specify only incremental changes to the router's labeled routing tree instead of the complete labeled routing tree.

The new router's labeled routing tree $newLRT$ is compared to the last reported labeled routing tree ($\{LRT_i\}$ for LORA and LRT_i for ORA) (procedure `ReportChanges`), and an update message that will be broadcast to the neighbors is constructed from the differences of the two trees. An RSU is generated if the link is in the new labeled routing tree but not in the current labeled routing tree, or if the parameters of the link have changed. For the case of a router running LORA, the labeled routing trees are only compared with each other if at least one of the four rules described in section V, below, is met, i.e., $M_i = \text{TRUE}$.

If the new router's labeled routing tree was compared against the last reported labeled routing tree then the router removes from the topology graph all the links that

are no longer used by any neighbor in their labeled routing trees (failed links are only removed from the topology graph by aging).

Finally, the current shortest-path tree LRT_k is discarded and the new one becomes the current labeled routing tree. The router's labeled routing tree is then used to compute the new routing table, using for example a depth-first search in the shortest-path tree.

II. Information Exchanged in AIR

Prior routing protocols based on topology information or distance information are based on the parameters of links exclusively. In contrast, AIR uses an update unit that conveys information about the performance characteristics and addressing information for a link and the node at the end of the link. More specifically, an RSU includes the following elements:

- a) A time stamp that validates the RSU;
- b) A type-of-service vector;
- c) The network address of the head node of the link;
- d) The network address of the tail node of the link;
- e) The link state parameters of the link between the two IRs; and
- f) The node state parameters of the tail of the link.

An update message sent by an IR contains at least one RSU. The time stamp of the RSU is assigned by the head of the link and should not be altered by any other node relaying the RSU in an update message. The type-of-service (TOS) vector is a bit vector specifying the TOS routing tree in which the link is being used by the node sending the RSU. The state parameters of a link are specified as a list of tuples, with each tuple consisting of a type and a content. There are two classes of state parameters for a link: performance parameters and addressing parameters. The performance of a link can be characterized in terms of its delay, cost, bandwidth, and reliability, for example. An addressing parameter specifies an identifier assigned to the link. An example of such an identifier in the present scheme is the local link identifier (LLID) assigned to the link by

the head of the link. The state parameters of the tail of a link may include, for example, the remaining battery life of the node.

III. Information Stored in AIR

Figures 2 a - 2d illustrate the fact that IRs running AIR need maintain only partial topology information. These illustrations provide an example of a six-node wireless network (each node labeled a - f, respectively). For simplicity, these Figures assume that a single link parameter is used to characterize a link in one of its directions, called the "cost" of the directed link. For example, link 32c may have a cost of 5 in the direction b to c, but only a cost of 1 in the direction c to b. In other examples, nodes may be coupled by multiple links between them, each having the same or different costs. Figures 2b through 2d show the selected topology according to AIR at the IRs indicated with filled circles. Solid lines represent the links that are part of the labeled routing tree of the respective IR. Arrowheads on links indicate the direction of the link stored in the IR's topology graph. IR a's labeled routing tree shown in Figure 2b is formed by the labeled routing trees reported by its neighbors b and c, and the links for which IR a is the head node (namely links (a, b) and (a, c)). Similarly, Figure 2c shows the LRT for IR b and Figure 2d that for IR c. From the figures, the savings in storage requirements should be clear, even for this few-node network.

The information maintained by an IR to participate in AIR includes a topology graph, an LRT for each TOS defined in the network, a routing table, and an adjacent-link table. The record entry for the link from u to v in the topology graph consists of the tuple (u, v, t, {l}, {n}), where u and v are the network addresses of the head and tail of the link, respectively, t is the most recent time stamp received for link (u, v), {l} is a sequence of type-value pairs specifying link parameters, and {n} is a sequence of type-value pairs specifying node parameters. A link parameter used in the present scheme is the LLID of the link.

The routing table specifies, for each destination and for each TOS, the next IR in the path to the destination and the distance to that destination based on the distance metric used for the TOS.

The cost of a failed link is considered to be infinity for any TOS. There are various ways in which costs may be assigned to links for a given TOS known in the art. For example, the cost of a link could simply be the number of hops, or the addition of the latency over the link plus some constant bias.

IV. Validating Updates

Because of delays in the IRs and links of an internetwork, update messages sent by a IR may propagate at different speeds along different paths. Therefore, a given IR may receive an RSU from a neighbor with stale link-state information, and a distributed termination-detection mechanism is necessary for a IR to ascertain when a given RSU is valid and avoid the possibility of RSUs circulating forever. AIR uses time stamp to validate RSUs. An IR either maintains a clock that does not reset when the IR stops operating, or asks its neighbors for the oldest known time stamp after it initializes or reboots.

An IR receiving an RSU accepts the RSU as valid if the received RSU has a larger time stamp than the time stamp of the RSU stored from the same source, or if there is no entry for the link in the topology graph and the RSU is not reporting an infinite cost. Link-state information for failed links are the only RSUs erased from the topology graph due to aging (which may be on the order of an hour or so (or other time period) after having processed the RSU). RSUs for operational links are erased from the topology graph when the links are erased from the routing trees of all the neighbors.

It is noted that, because RSUs for operational links never age out, there is no need for the head node of a link to send periodic RSUs to update the time stamp of the link. This is important, because it means that AIR does not need periodic update messages to validate link-state information like OSPF, J. Moy, "OSPF Version 2," RFC

1583, Network Working Group, March 1994, and all prior protocols based on sequence numbers or time stamps.

V. Exchanging Update Messages

An IR sends RSUs in two different ways: (a) Following an optimum routing approach; and (b) following a least overhead approach. The optimum routing approach is well suited for networks with fairly static topologies. The least overhead approach is tailored for networks with dynamic topologies due to IR mobility. Which approach should be executed in IRs can be defined by an IR configuration parameter.

As indicated in the pseudocode description, according to the optimum routing approach an IR sends RSUs about a link in the following cases: (a) when an RSU is received for the link causing the link to be added to the IR's routing tree, (b) when the link is already in the IR's routing tree and a more recent RSU is received for the link, (c) when an RSU reporting the failure of the link results in no path to the tail of the link, and (d) when a failed link is not in the IR's routing tree and there is no path to the tail of the link.

In contrast, according to the least overhead approach an IR sends RSUs according to the following rules:

- (a) The IR finds a new destination, or any of its neighbors reports a new destination.
- (b) The IR finds that the change in the cost of a path to at least one destination exceeds a threshold delta, or at least one destination becomes unreachable to the IR or any of its neighbors.
- (c) A path implied in the LRT of the IR leads to a loop.
- (d) The IR sends an RSU when: (i) The new successor chosen to a given destination has an address larger than the address of the IR; and (ii) the reported distance from the new chosen successor n to a destination j is longer than the reported distance from the previous

successor to the same destination. However, if the link from the IR to j fails and n is a neighbor of j , no update message is needed regarding j or any destination whose path from the IR involves j .

Each time an IR processes an update message from a neighbor, it updates that neighbor's LRT and traverses that tree to determine for which destinations its neighbor uses the IR as a relay in its preferred paths. The IR then determines if it is using the same neighbor as a relay for any of the same destinations. A routing loop is detected if the IR and neighbor use each other as relay to any destination, in which case the loop must be broken and the IR must send an update message with the corresponding changes.

We observe that, in any routing loop among IRs with unique addresses, one of the IRs must have the smallest address in the loop; therefore, if an IR is forced to send an update message when it chooses a successor whose address is larger than its own, then it is not possible for all IRs in a routing loop to remain quiet after choosing one another, because at least one of them is forced to send an update message, which causes the loop to break when IRs update their LRTs.

To ensure that AIR works correctly with least overhead routing and incremental updates specifying only changes to an LRT, an IR must remember the LRT that was last notified to its neighbors. If any of the rules for update notification in least overhead routing are satisfied, the IR must do one of two things: (a) If the LRT includes new neighbors than those present in the LRT that was last updated, then the IR must send its entire LRT in its update, so that new neighbors learn about all the destinations the IR knows; (b) if the two LRTs imply the same neighbors, the IR sends only the updates needed to obtain the new LRT from the old one.

To ensure that AIR stops sending update messages, a simple rule can be used to determine which IR must stop using its neighbor as a relay, such a rule can be, for example, "the IR with the smaller address must change its path."

The rules for update-message exchange according to least overhead routing stated above assume that an update message is sent reliably to all the neighbors of an IR.

The following example illustrates a scenario in which the last rule is needed to prevent permanent loops. Consider the six-node wireless network 30 shown in Figure 2a. In this example, IRs are given identifiers that are lexicographically ordered, i.e., "a" is the smallest identifier and "f" is the largest identifier in the graph. All links and nodes are assumed to have the same propagation delays, all the links but links (a, b) and (b, c) have unit cost, and $\delta = \infty$. Figures 2b through 2d show the LRTs according to AIR at the IRs indicated with filled circles for the network topology depicted in Figure 2a. Arrowheads on solid lines indicate the direction of the links stored in the IR's LRT.

Figures 3a - 3f now depict the sequence of events triggered by the execution of AIR in the example network after the failures of links (c, d) 32g and (b, e) 32d. The figures show the RSUs (in parentheses) generated by the node with filled circle, which RSUs are transmitted in an update message to the node's neighbors. The third element in an RSU corresponds to the cost of the link (a RESET has cost infinity). As shown in Figure 3b, node c transmits an RSU after processing the failure of link (c, d) 32g; the distance from the new successor b to d and f is longer than from the previous successor d. When link (b, e) 32d fails (see Figure 3d), node b realizes that the destinations d, e, and f are unreachable and generates an RSU reporting the failure of the link connecting to the head of the subtree of the LRT that becomes unreachable. The RSU from b triggers the RSUs that allow nodes a, b, and c to realize that there are no paths to d, e, and f (Figures 3e and 3f). A similar sequence of events takes place at the other side of the network partition (not shown).

As another example of the operation of AIR, consider the seven-node wireless network 40 shown in Figure 4a. All links and nodes are assumed to have the same propagation delays, all the links have unit cost, $\delta = \infty$. Figures 4b through 4d show the LRTs produced according to AIR at the IRs indicated with filled circles for the network topology depicted in Figure 4a. Arrowheads on solid lines indicate the direction of the links stored in the IR's LRT. When the link (f, g) fails (Figure 4e), the neighbor protocol at node f triggers the execution of procedure NeighborDown, the link (d, g) is inserted into f's LRT but no update message is generated because f's new successor

towards g has an address smaller than f and destination g is a neighbor of the new successor. Figure 4f shows the new LRT of node d after the failure of link (d, g). Because d has an address smaller than the new successor towards g it is required to send an update message reporting the new link added to the LRT. Nodes c, e, and f do not generate any update message after processing d's message because there exist a path to all destinations in the network and no routing loop was formed. This example thus illustrates how link failures may not cause the generation of update messages by nodes that have the failed link in their LRTs as long as the nodes have a path to all destinations.

VI. Obtaining LRTs and Source Routes

Obtaining the LRTs and source routes in the present scheme is done with a very simple modification to Dijkstra's SPF algorithm that is run on the topology graph of an IR. The modifications to SPF consist of checking for the proper bit to be set in the TOS bit vector of a link so that only those links being used for the required TOS are considered in preferred paths, and accumulating the source route in terms of LLIDs as the topology graph is traversed.

Given the topology graph, the algorithm proceeds as follows:

1. Initialize:

Value of TOS bit vector = T

Set IR-set = {root}, where root is the IR running the algorithm.

Dist-to-root = 0

Route-to-root = root

for all IRs other than root set

Dist-to-IR = infinity

Route-to-IR = root

for all IRs neighbors of root

If TOS bit vector = T then

Dist-to-IR = cost of link to IR

Route-to-IR = Route-to-IR \cup LLID of link to IR

2. Find next IR in IR-set:

Find an IR x not in set such that:

Dist-to-IR = Minimum{ Dist-to-IR not in IR-set

considering only links with TOS bit vector = T }

Augment IR-set = IR-set \cup x

Stop if IR-set contains all IRs

3. Change chosen distance and path:

For each IR y not in IR-set do:

Dist-to-y = Minimum{ Dist-to-y, Dist-to-z + cost of link(z,y)

with z in IR-set and TOS bit vector of (z,y) = T }

Route-to-y = Route-to-z \cup LLID of (z,y)

4. Repeat Step 2.

Thus a scheme for enabling routing of data packets in a computer network has been described. Although the foregoing description and accompanying figures discuss and illustrate specific embodiments, it should be appreciated that the present invention is to be measured only in terms of the claims that follow the example of a pseudocode listing for the routing protocol set forth below:

CLAIMS

What is claimed is:

1. A method, comprising producing one or more labeled routing trees (LRTs) at a router of a computer network according to a shortest path determination made over a partial topology graph of the network, which graph is produced according to knowledge of adjacent links of the router and one or more LRTs of neighboring routers.
2. The method of claim 1 wherein the LRTs of the router are updated in response to receipt of routing state update messages.
3. The method of claim 2 wherein the routing state update messages include local link identifiers assigned by a head of a link to which the identifiers pertain.
4. The method of claim 3 wherein the routing state update messages further include node parameters of a tail of the link to which the local link identifiers pertain.
5. The method of claim 3 wherein the routing state update messages are transmitted within the network: (i) in response to a new destination node being detected by an existing node within the network, (ii) in response to a destination becoming unreachable by a collection of the existing nodes, and/or (iii) in response to the change in the cost of a path to at least one destination exceeding a threshold delta, (iv) in situations where a routing loop may be encountered among two or more of the nodes of the network.
6. The method of claim 5 wherein situations where a routing loop may be encountered include times when a path implied in the LRT of the router leads to a loop.
7. A routing protocol, comprising distributing local link identifiers among nodes of a computer network within routing state update messages.
8. The routing protocol of claim 7 wherein the local link identifiers are each assigned by a head of a link to which the identifiers pertain.
9. The routing protocol of claim 7 wherein the routing state update messages include state parameters for nodes at tails of links to which the identifiers pertain.

10. The routing protocol of claim 7, further comprising distributing labeled routing trees of the nodes of the computer network among the nodes.

11. The routing protocol of claim 10 wherein each node of the computers network distributes its labeled routing trees to its neighboring nodes.

12. The routing protocol of claim 11 wherein each node of the computer network maintains one labeled routing tree per type of service offered in the network.

13. The routing protocol of claim 11 wherein each node generates its labeled routing trees according to a partial topology graph derived from the node's adjacent links and labeled routing trees of neighboring nodes.

14. The routing protocol of claim 13 wherein the labeled routing trees of a node are generated by applying a path selection algorithm to the partial topology graph.

15. The routing protocol of claim 14 wherein each node of the computer network maintains one labeled routing tree per type of service offered in the network.

16. The routing protocol of claim 15 wherein labeled routing trees of nodes in the computer network are updated in response to receipt of the routing state update messages.

17. The routing protocol of claim 16 wherein labeled routing trees are updated according to whether an optimum routing approach or a least overhead routing approach is used within the computer network.

18. The routing protocol of claim 17 wherein in the optimum routing approach, routing state update messages are transmitted: (i) when a routing state update message which causes a link to be added to a subject node's labeled routing trees is received, (ii) when a more recent routing state update message regarding an existing link in the subject node's labeled routing trees is received, and/or (iii) when a destination becomes unreachable by the subject node.

19. The routing protocol of claim 17 wherein in the least overhead routing approach, routing state update messages are transmitted: (i) when a subject node finds a new

destination or any neighbors of the subject node report same, (ii) when a destination becomes unreachable by the subject node or any of its neighbors, (iii) when the subject node finds that the change in the cost of path to at least one destination exceeds a threshold delta, (iv) when a path implied in any of the labeled routing trees of the subject node leads to a loop, and/or a new successor chosen to a given destination has an address larger than an address of the subject node and a reported distance from the new successor to a particular destination is larger than a reported distance from a previous successor to that particular destination.

20. A routing update message, comprising information regarding performance characteristics and addressing information for a link of a computer network and a node of the network at a tail end of the link.

21. The routing update message of claim 20 wherein the routing update message further comprises a time stamp assigned by a node at a head end of the link.

22. The routing update message of claim 21 wherein the performance characteristics include link state parameters of the link.

23. The routing update message of claim 21 wherein the routing update message further comprises a type of service vector.

24. The routing update message of claim 21 wherein the routing update message further comprises addressing information for a node at a head end of the link.

25. The routing update message of claim 23 wherein the type of service vector specifies a type of service routing tree in which the link is being used by a node transmitting the routing update message.

26. The routing update message of claim 20 wherein the performance characteristics of the link comprise link delay, link cost, available bandwidth, and/or reliability.

27. The routing update message of claim 20 wherein the addressing information of the link comprises a local link identifier assigned to the link by a node at a head end of the link.

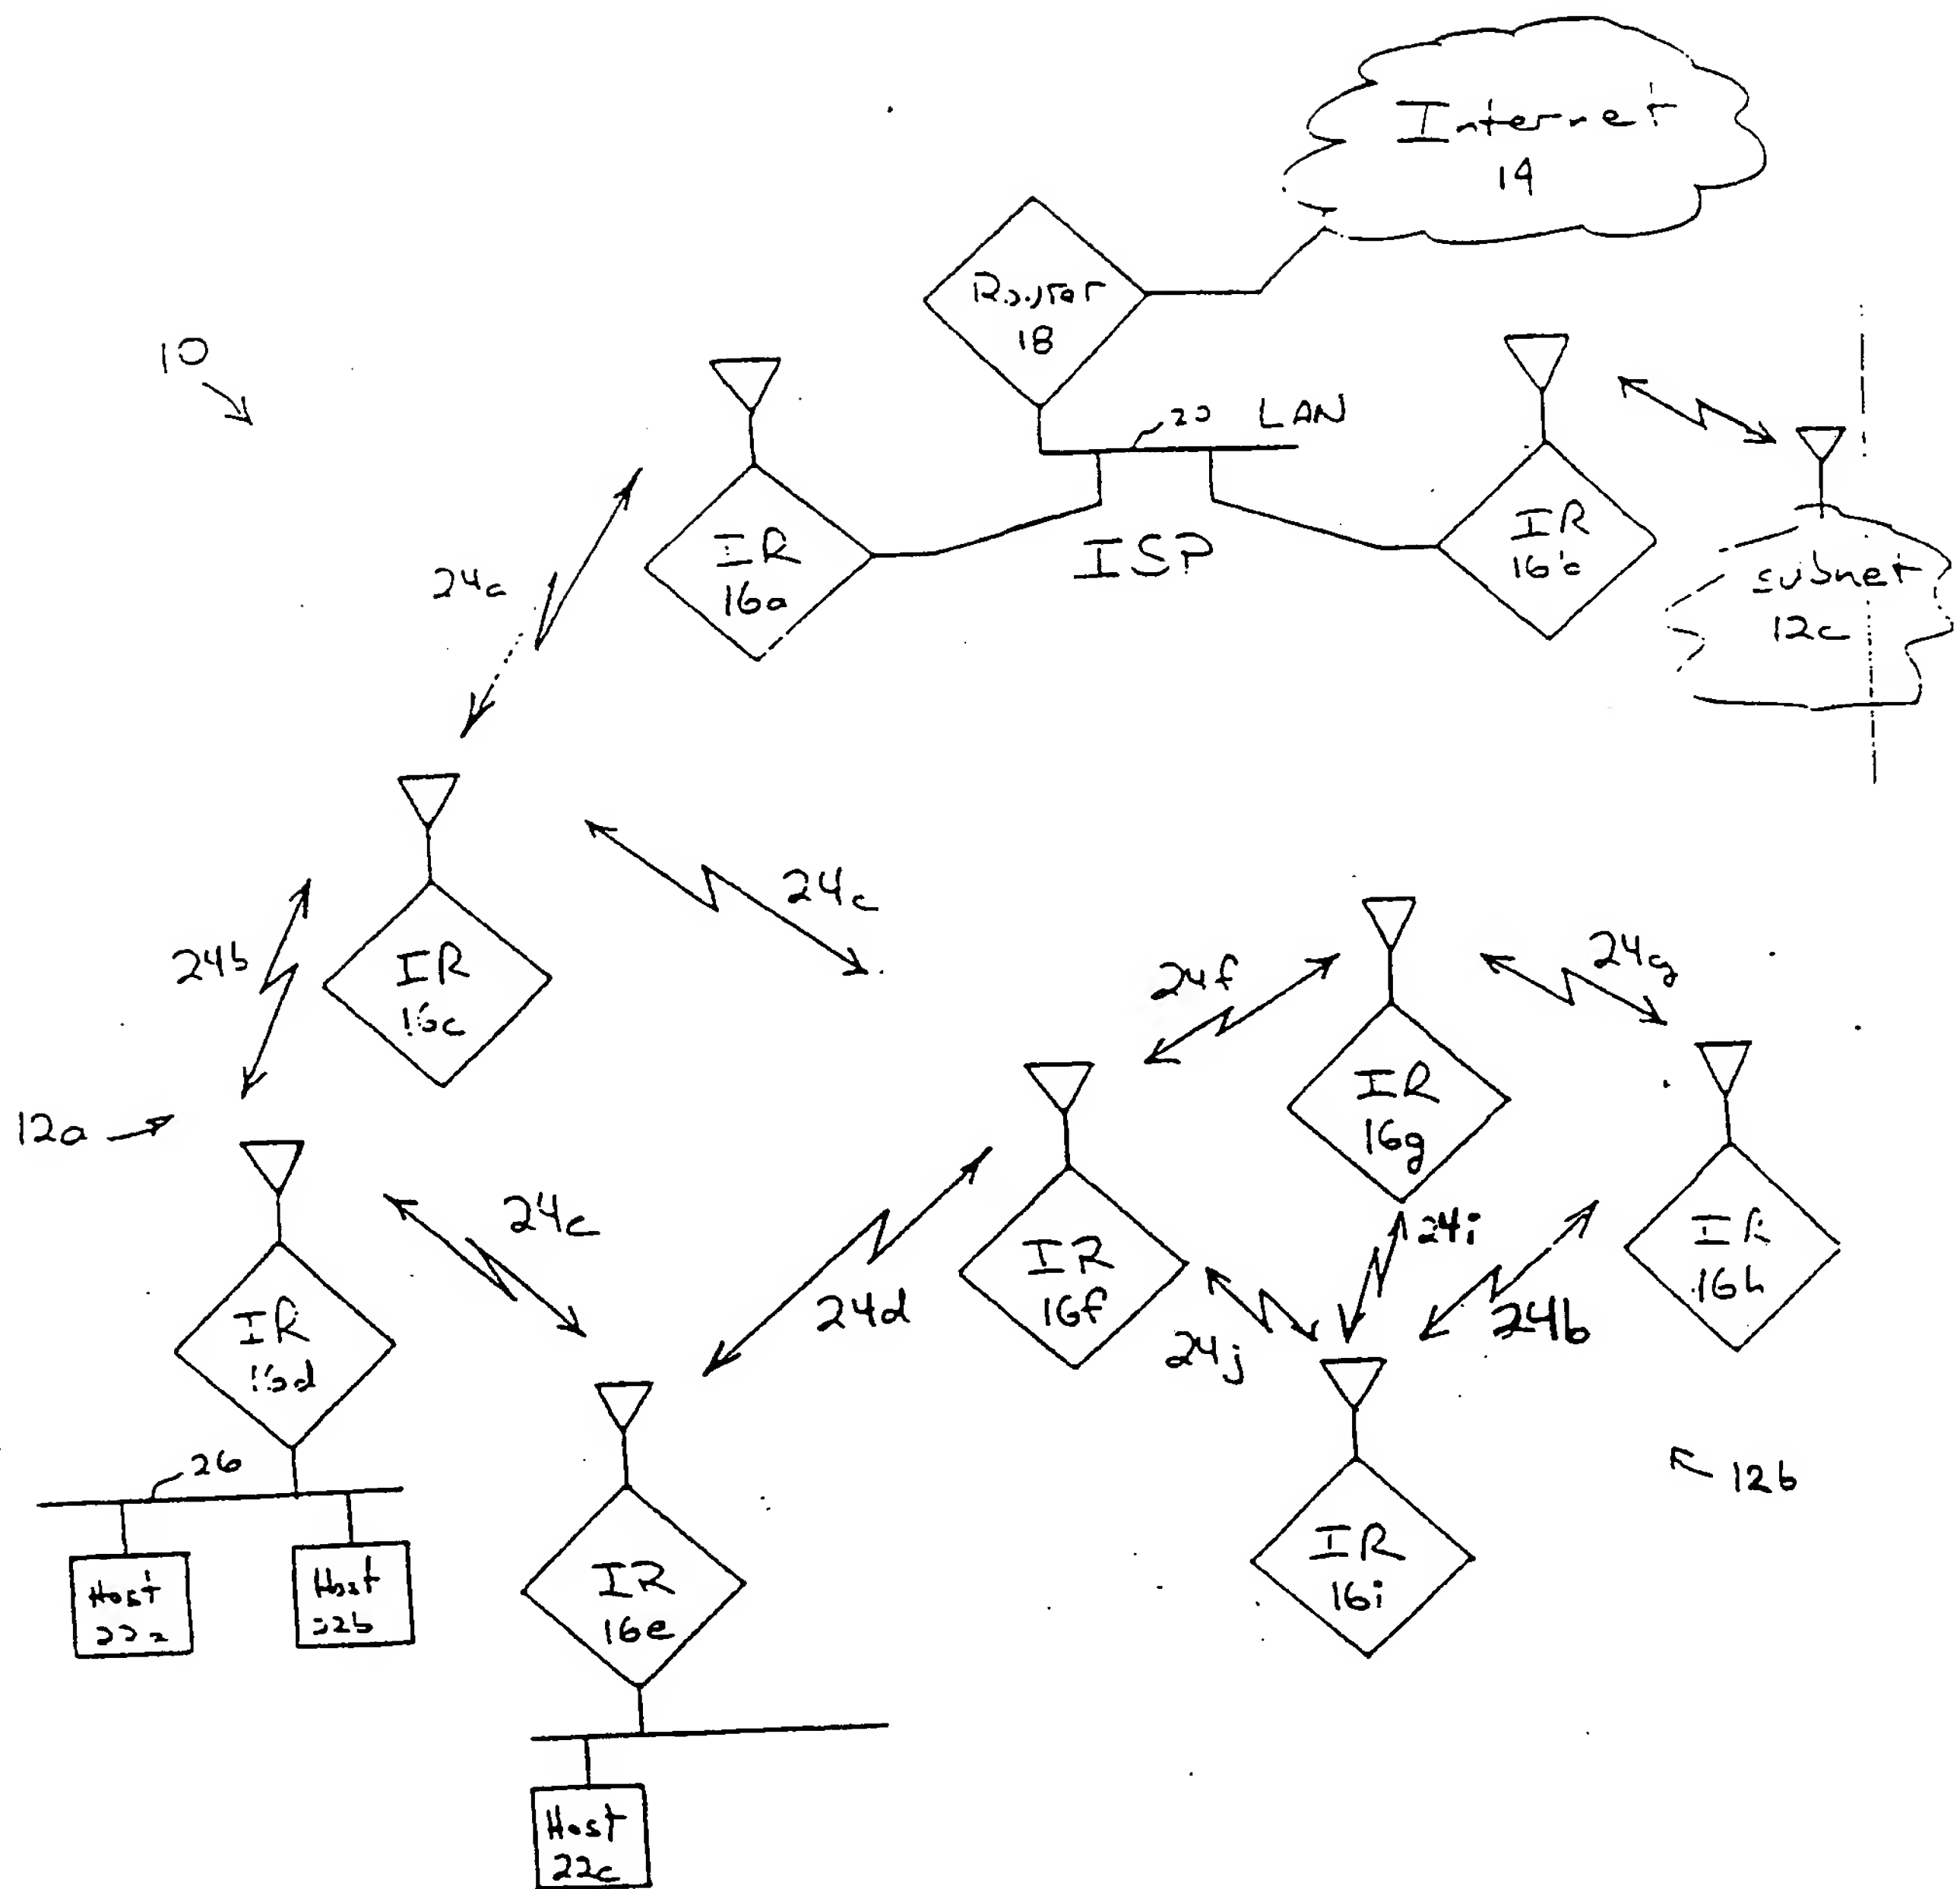


FIG. 1.

Fig. 2a




Fig. 2b




Fig. 2c





Fig. 2d

Fig. 3a




Fig. 3b




Fig. 3c





Fig. 3d




Fig. 3e




Fig. 3f

Fig. 4a




Fig. 4b




Fig. 4c





Fig. 4d




Fig. 4e




Fig. 4f

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
26 April 2001 (26.04.2001)

PCT

(10) International Publication Number
WO 01/30035 A3

- (51) International Patent Classification⁷: **H04L 12/56**
- (21) International Application Number: **PCT/US00/41180**
- (22) International Filing Date: **16 October 2000 (16.10.2000)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:
09/418,700 15 October 1999 (15.10.1999) **US**
- (71) Applicant (for all designated States except US): **NOKIA WIRELESS ROUTERS, INC.** [US/US]: 313 Fairchild Drive, Mountain View, CA 94043 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **GARCIA-LUNA-ACEVES, J. Joaquin** [MX/US]: 82 Lakewood Circle, San Mateo, CA 94402 (US). **SPOHN, Marcelo** [BR/US]:

623 Kushland Way, Santa Cruz, CA 95064 (US). **BEYER, David, A.** [US/US]: 468 Paco Drive, Los Altos, CA 94024 (US).

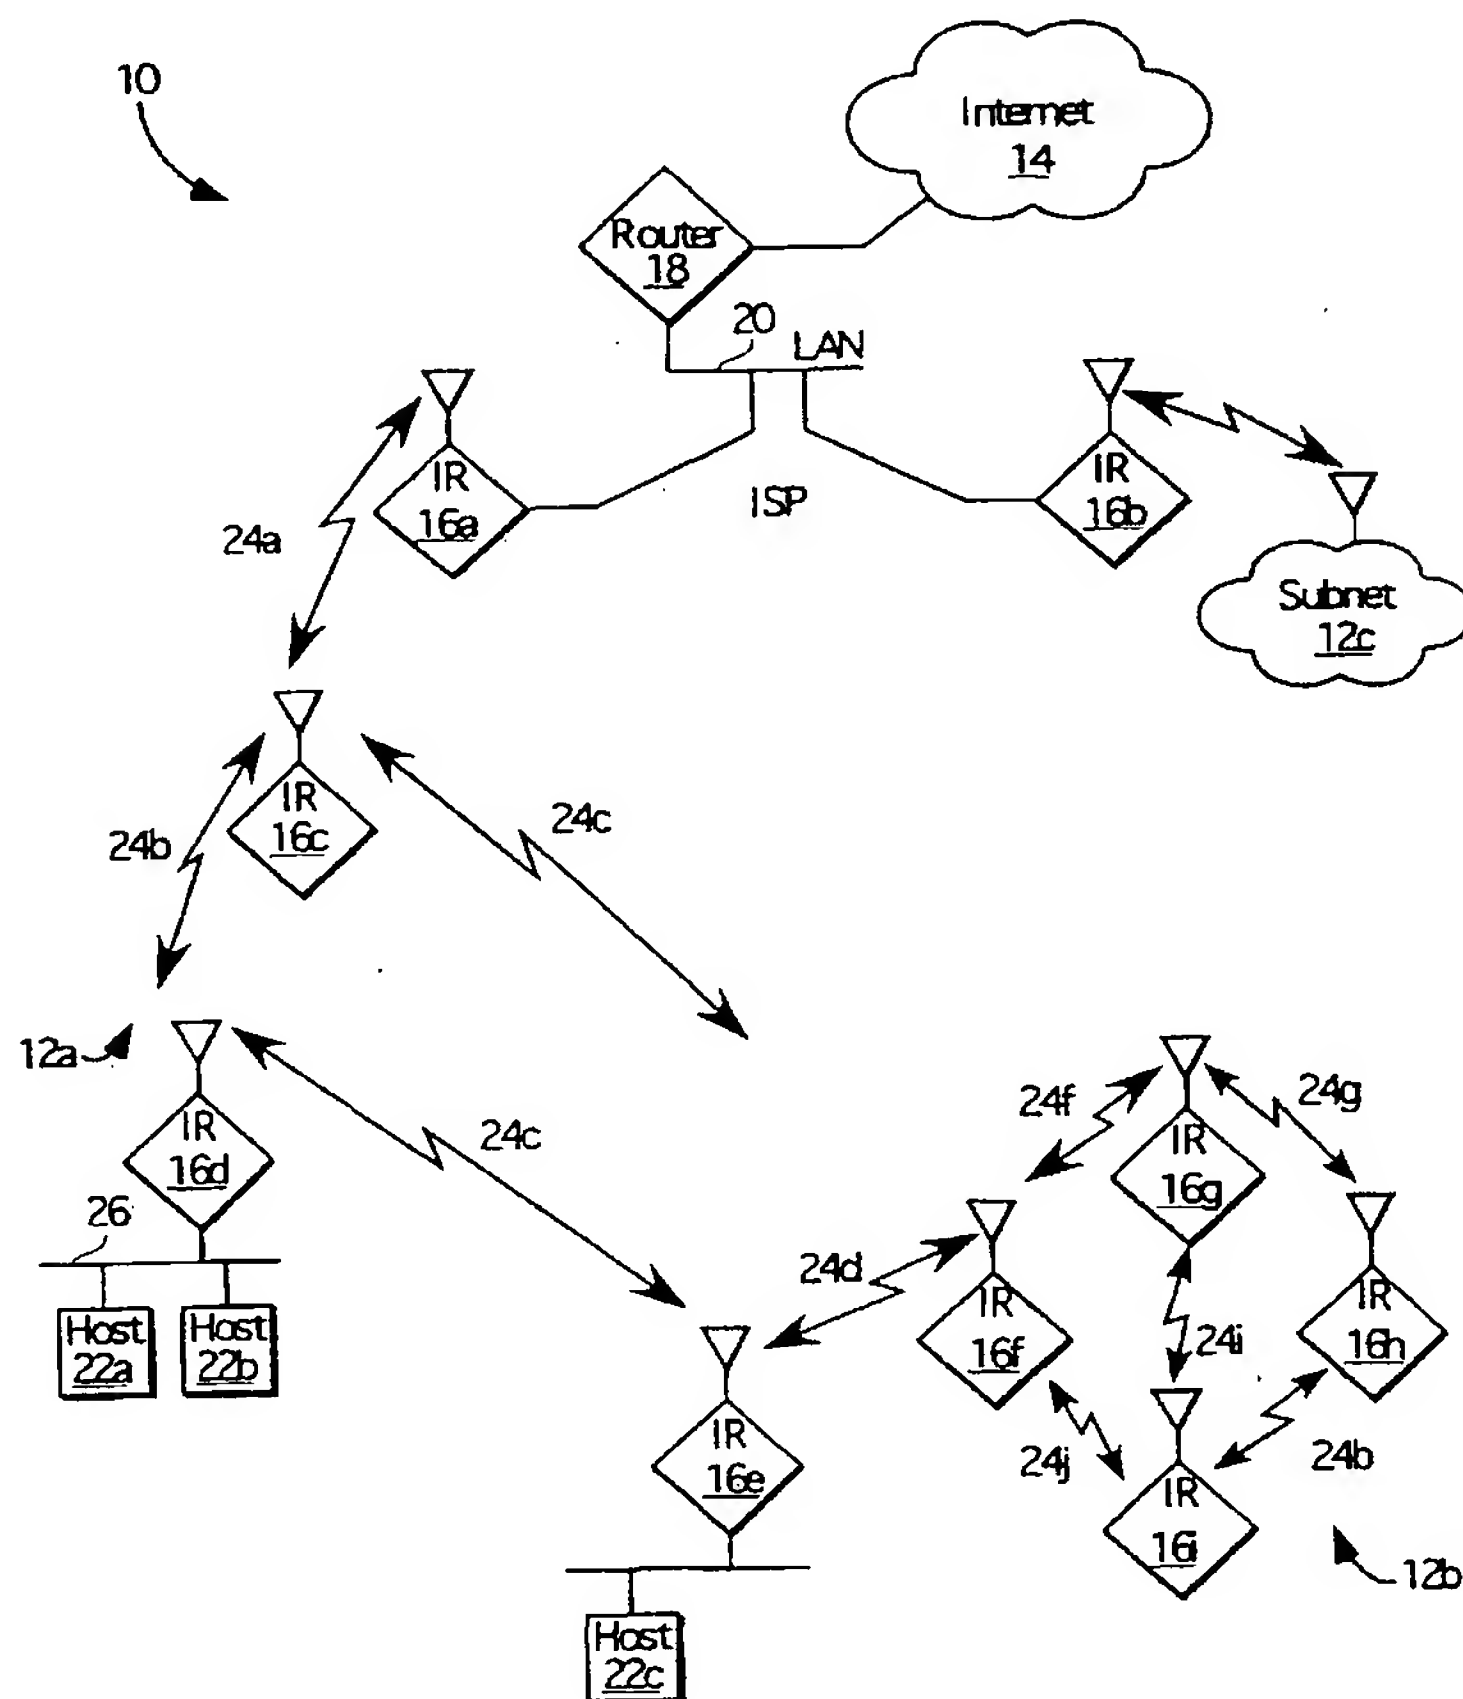
(74) Agents: **FAHMI, Tarek, N.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 12400 Wilshire Boulevard, 7th floor, Los Angeles, CA 90025 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,

[Continued on next page]

(54) Title: **SYSTEM FOR COMMUNICATING LABELED ROUTING TREES**



(57) Abstract: One or more labeled routing trees (LRTs) are produced at a router of a computer network according to a shortest path determination made over a partial topology graph of the network, which graph is produced according to knowledge of adjacent links of the router and or more LRTs of neighboring routers. The LRTs of the router may be updated in response to receipt of routing state update messages, and such messages may include local link identifiers assigned by a head of a link to which the identifiers pertain, and node parameters of a tail of the link to which the local link identifiers pertain. The routing state update messages may be transmitted within the network: (i) in response to a new destination node being detected by an existing node within the network, (ii) in response to a destination becoming unreachable by a collection of the existing nodes, (iii) in response to the change in the cost of a path to at least one destination exceeding a threshold and/or (iv) in situations where a routing loop may be encountered among two or more of the nodes of the network (e.g., at times when a path implied in the LRT of the router leads to a loop).



IT, LU, MC, NL, PT, SE). OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

(88) Date of publication of the international search report:
6 December 2001

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/41180

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WCNC. 1999 IEEE WIRELESS COMMUNICATIONS AND NETWORKING CONFERENCE, vol. 03, 24 September 1999 (1999-09-24), pages 1308-1312, XP002168435 Piscataway, NJ, USA, IEEE, USA ISBN: 0-7803-5668-3 IEEE Catalog Number: 99TH8466 page 1308, right-hand column, line 10 -page 1309, left-hand column, line 38 -----	1-27

☐ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

*** Special categories of cited documents:**

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

29 May 2001

Date of mailing of the international search report

15/06/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Siebel, C

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
26 April 2001 (26.04.2001)

PCT

(10) International Publication Number
WO 01/030035 A3

(51) International Patent Classification⁷: **H04L 12/56**

(21) International Application Number: **PCT/US00/41180**

(22) International Filing Date: **16 October 2000 (16.10.2000)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
09/418,700 15 October 1999 (15.10.1999) **US**

(71) Applicant (for all designated States except US): **NOKIA WIRELESS ROUTERS, INC.** [US/US]; 313 Fairchild Drive, Mountain View, CA 94043 (US).

(72) Inventors; and

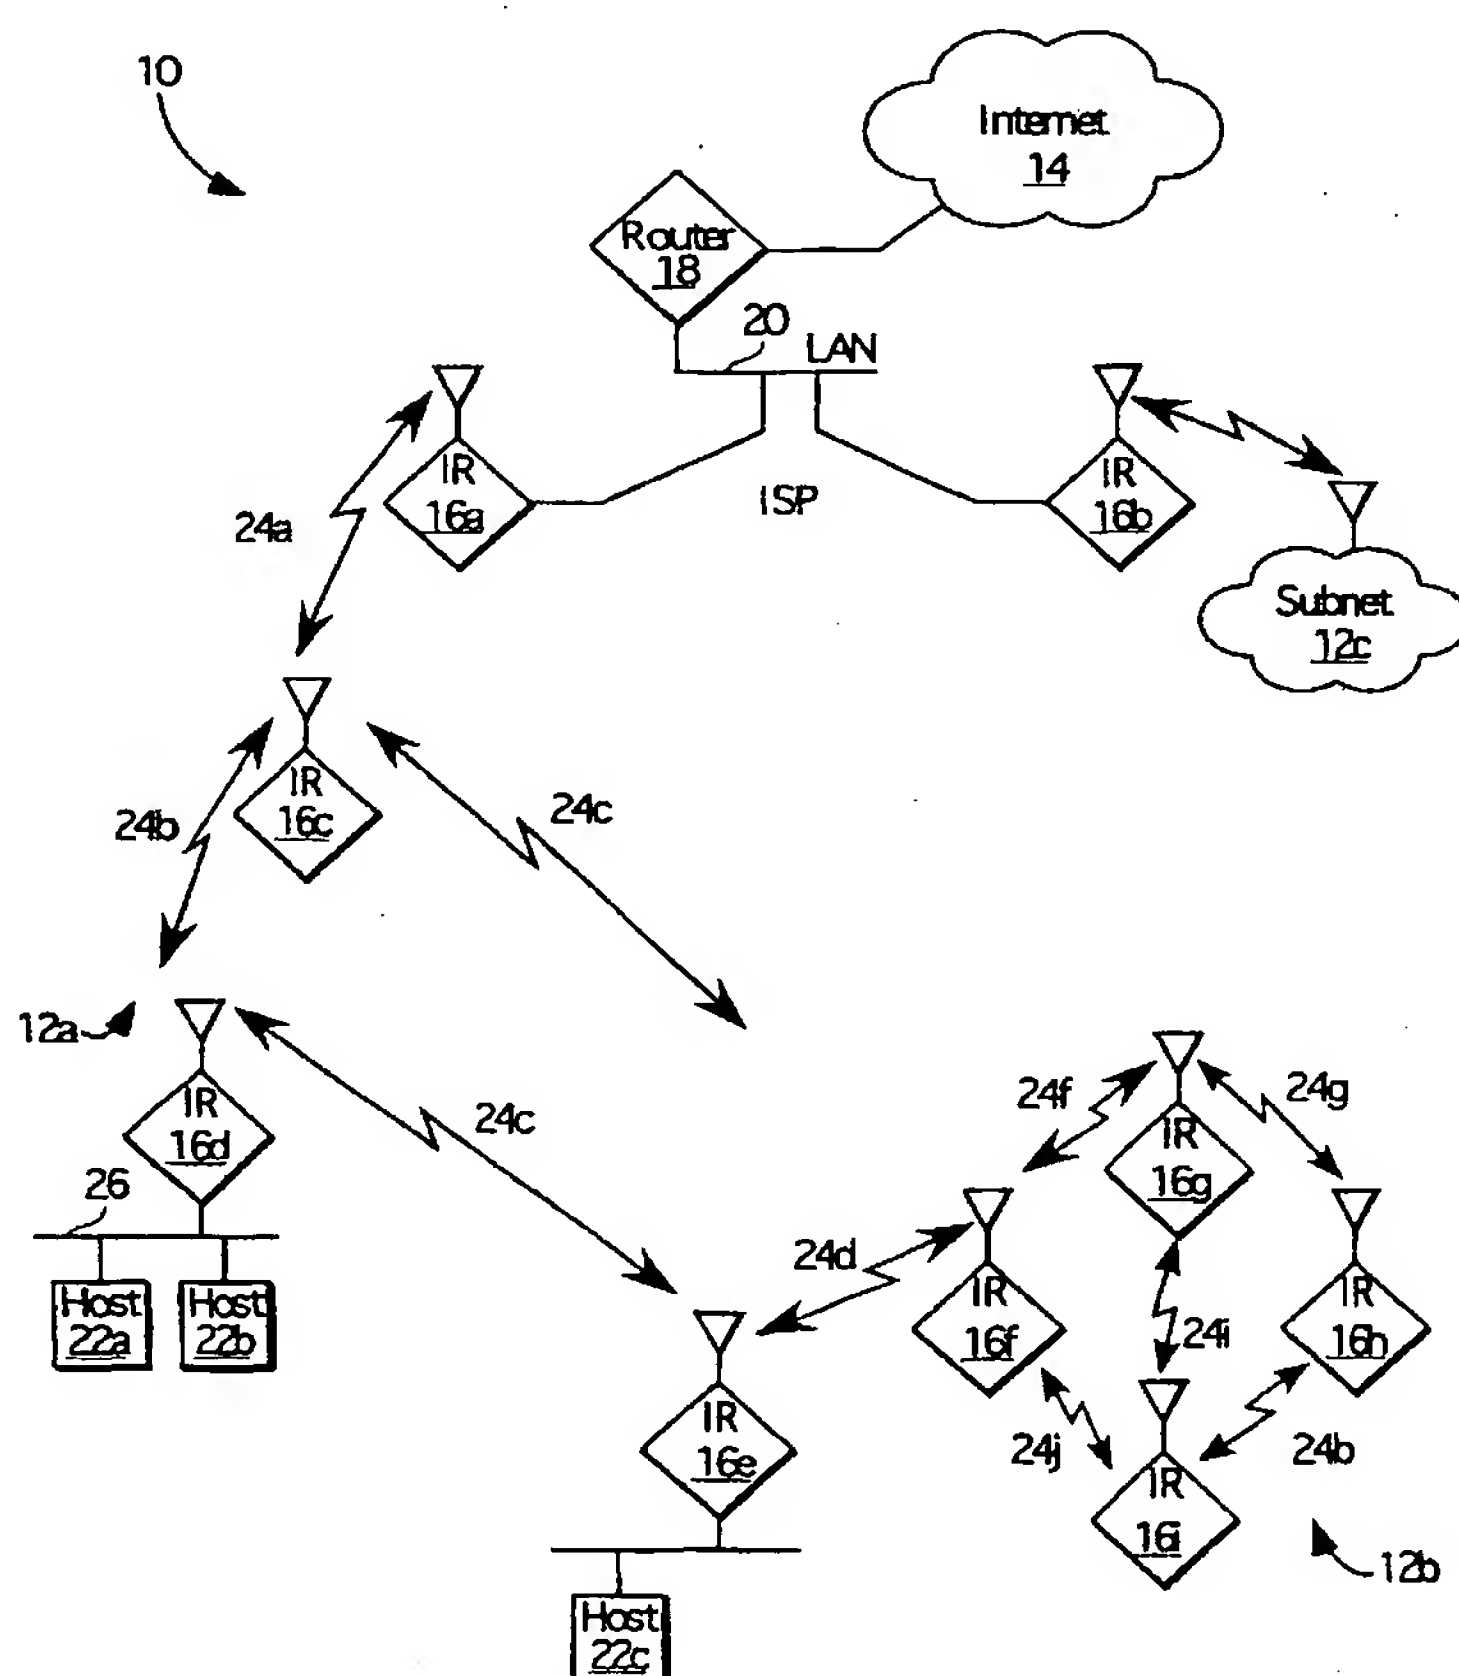
(75) Inventors/Applicants (for US only): **GARCIA-LUNA-ACEVES, J. Joaquin** [MX/US]; 82 Lakewood Circle, San Mateo, CA 94402 (US). **SPOHN, Marcelo** [BR/US]; 623 Kushland Way, Santa Cruz, CA 95064 (US). **BEYER, David, A.** [US/US]; 468 Paco Drive, Los Altos, CA 94024 (US).

(74) Agents: **FAHMI, Tarek, N.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 12400 Wilshire Boulevard, 7th floor, Los Angeles, CA 90025 (US).

(81) Designated States (national): **AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ,**

[Continued on next page]

(54) Title: **SYSTEM FOR COMMUNICATING LABELED ROUTING TREES**



(57) Abstract: One or more labeled routing trees (LRTs) are produced at a router of a computer network according to a shortest path determination made over a partial topology graph of the network, which graph is produced according to knowledge of adjacent links of the router and or more LRTs of neighboring routers. The LRTs of the router may be updated in response to receipt of routing state update messages, and such messages may include local link identifiers assigned by a head of a link to which the identifiers pertain, and node parameters of a tail of the link to which the local link identifiers pertain. The routing state update messages may be transmitted within the network: (i) in response to a new destination node being detected by an existing node within the network, (ii) in response to a destination becoming unreachable by a collection of the existing nodes, (iii) in response to the change in the cost of a path to at least one destination exceeding a threshold and/or (iv) in situations where a routing loop may be encountered among two or more of the nodes of the network (e.g., at times when a path implied in the LRT of the router leads to a loop).

WO 01/030035 A3



NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(88) Date of publication of the international search report:
6 December 2001

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

(48) Date of publication of this corrected version:
1 August 2002

(15) Information about Correction:
see PCT Gazette No. 31/2002 of 1 August 2002, Section II

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM FOR COMMUNICATING LABELED ROUTING TREES

STATEMENT OF GOVERNMENT LICENSE RIGHTS

The United States Government has a paid-up license in portions of this invention and the right in limited circumstances to require the patent owner to license others on reasonable terms as provided for by the terms of Contract No.: DAAH01-98-C-R005, awarded by the U.S. Army Aviation & Missile Command.

FIELD OF THE INVENTION

The present invention relates to routing protocols in computer networks, and more particularly, routing protocols in ad hoc networks with radio links.

BACKGROUND

Multi-hop packet-radio networks, or ad hoc networks, consist of mobile hosts interconnected by routers that can also move. The deployment of such routers is ad hoc and the topology of such networks is very dynamic, because of host and router mobility, signal loss and interference, and power outages. In addition, the channel bandwidth available in ad hoc networks is relatively limited compared to wired networks, and untethered routers may need to operate with battery-life constraints. In these networks, routing must preferably be accomplished using as few a number of control messages and neighbor-to-neighbor handshakes as possible, in order to conserve channel bandwidth for user data and preserve the battery life of untethered nodes. Because of the dynamics of the topology in an ad hoc network, broadcast radio links are preferable for interconnecting routers without the need for topology planning.

Routing protocols for computer networks can be categorized according to: (a) the type of information the protocols use to compute preferred paths, and (b) the way in which routers obtain routing information. In terms of the type of information used by routing protocols, routing protocols can be classified into link-state protocols and

distance-vector protocols. Routers running a link-state protocol use topology information to make routing decisions; routers running a distance-vector protocol use distances and, in some cases, path information, to destinations to make routing decisions.

In terms of the way in which routers obtain information, routing protocols have been classified as either table-driven or on-demand. In an on-demand routing protocol, routers maintain path information for only those destinations that they need to contact as a source or relay of information. The basic approach consists of allowing a router that does not know how to reach a destination to send a flood-search message to obtain the path information it needs. One of the first routing protocols of this type was proposed to establish virtual circuits in the MSE network, see V.O.K. Li and R. Chang, "Proposed routing algorithms for the US Army mobile subscriber equipment (MSE) network," Proc. IEEE MILCOM'86, Monterey, California, October 1986, and there are several other, recent examples of this approach (e.g., AODV, see C. Perkins, "Ad Hoc On Demand Distance Vector (AODV) Routing," draft-ietf-manet-aodv-00.txt, November 1997; ABR, see C-K. Toh, "Wireless ATM & Ad Hoc Networks," Kluwer, November 1996; DSR, see D. Johnson and D. Maltz, "Protocols for Adaptive Wireless and Mobile Networking," IEEE Pers. Commun., Vol. 3, No. 1, February 1996; TORA, see V. Park and M. Corson, "A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks," Proc. IEEE INFOCOM 97, Kobe, Japan, April 1997; SSA, see R. Dube et al., "Signal Stability-Based Adaptive Routing (SSA) for Ad Hoc Mobile Networks," IEEE Pers. Commun., February 1997; and ZRP, see Z. Haas and M. Pearlman, "The Zone Routing Protocol for Highly Reconfigurable Ad Hoc Networks," Proc. ACM SIGCOMM 98, Vancouver, British Columbia, August 1998). The Dynamic Source Routing (DSR) protocol has been shown to outperform many other on-demand routing protocols. J. Broch et al, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," Proc. ACM MOBICOM 98, Dallas, TX, October 1998. The existing on-demand routing protocols differ on the specific mechanisms used

to disseminate flood-search packets and the responses thereto, the means used to cache information received during other nodes' searches, and the manner in which to determine the cost of a link and the existence of a neighbor. However, one common characteristic of all of the on-demand routing protocols reported to date is that such protocols are based on distances to destinations. Stated differently, there have been no on-demand link-state routing protocol proposals to date.

In a table-driven scheme, each router maintains path information for each known destination in the network and updates its routing-table entries as needed. Examples of table-driven algorithms based on distance vectors are the routing protocol of the DARPA packet-radio network, J. Jubin and J. Tornow, "The DARPA Packet Radio Network Protocols," Proceedings of the IEEE, Vol. 75, No. 1, January 1987; DSDV, C. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers," Proc. ACM SIGCOMM 94, London, UK, October 1994; WRP, S. Murthy and J.J. Garcia-Luna-Aceves, "An Efficient Routing Protocol for Wireless Networks," ACM Mobile Networks and Applications Journal, Special issue on Routing in Mobile Communication Networks, 1996; WIRP, J.J. Garcia-Luna-Aceves et al., "Wireless Internet Gateways (WINGS)," Proc. IEEE MILCOM'97, Monterey, California, November 1997; and least-resistance routing protocols, M Pursley and H.B. Russell, "Routing in Frequency-Hop Packet Radio Networks with Partial-Band Jamming," IEEE Trans. Commun., Vol. 41, No. 7, pp. 1117-1124, 1993.

Prior table-driven approaches to link-state routing in packet-radio networks are based on topology broadcasts. However, disseminating complete link-state information to all routers incurs excessive communication overhead in an ad hoc network because of the dynamics of the network and the small bandwidth available. Accordingly, all existing link-state routing approaches for packet-radio networks have been based on hierarchical routing schemes. R. Ramanathan and M. Steenstrup, "Hierarchically-organized, Multihop Mobile Wireless Networks for Quality-of-Service Support," ACM Mobile Networks and Applications, Vol.3, No. 1, pp. 101-119, 1998; C.V.

Ramamoorthy and W. Tsai, "An Adaptive Hierarchical Routing Algorithm," Proceedings of IEEE COMPSAC '83, Chicago, Illinois, pp. 93-104, November 1983; and M. Steenstrup (ed.), Routing in Communication Networks, Prentice-Hall, 1995. Also, prior proposals for link-state routing using partial link-state data without clusters, see, e.g., J.J. Garcia-Luna-Aceves and J. Behrens, "Distributed, scalable routing based on vectors of link states," IEEE Journal on Selected Areas in Communications, Vol. 13, No. 8, 1995; and J.J. Garcia-Luna-Aceves and M. Spohn, "Scalable Link-State Internet Routing," Proc. IEEE International Conference on Network Protocols (ICNP 98), Austin, Texas, October 14-16, 1998, require routers to explicitly inform their neighbors which links they use and which links they stop using.

A number of prior routing protocols are based on the notion of routing trees, in which routers communicate either the state (i.e., cost or length) of the links in a shortest-path routing tree, or the distance from the root of the tree and the second-to-last hop in the routing tree for each node in the tree. An early example of this type of protocol was proposed in U.S. Patent 4,466,060 to Riddle. In Riddle's protocol, a router communicates different routing trees to different neighbors; such trees are called "exclusionary trees" and specify the preferred paths to destinations excluding those paths that involve the router to which the update is being sent. An update packet or message specifies an entire exclusionary tree. Another protocol based on routing trees was reported by Garcia-Luna-Aceves, J.J. Garcia-Luna-Aceves, "A Fail-Safe Routing Algorithm for Multihop Packet-Radio networks," Proc. IEEE Infocom 86, Miami, FL, April 1986, which protocol differs from Riddle's protocol in that the same routing tree is sent incrementally by a router to all its neighbors. A. Humblet, "Another Adaptive Shortest-Path Algorithm," IEEE Trans. Comm., Vol.39, No.6, June 1991, pp.995--1003 (and see US Patent 4,987,536); Cheng et al., C. Cheng et al., "A Loop-Free Extended Bellman-Ford Routing Protocol without Bouncing Effect", Proc. ACM SIGCOMM 89, pp.224--236; B. Rajagopalan and M. Faiman, "A Responsive Distributed Shortest-Path Routing Algorithm within Autonomous Systems," Journal of Internetworking: Research and Experience}, Vol. 2, No. 1, March 1991, pp. 51-69; and S. Murthy and J.J. Garcia-

Luna-Aceves, "Loop-Free Internet Routing Using Hierarchical Routing Trees," Proc. IEEE INFOCOM 97, Kobe, Japan, April 7-11, 1997, have all proposed protocols based on source trees in which a router communicates to its neighbors the same shortest-path routing tree incrementally and these latter examples differ from the above-cited protocol by Garcia-Luna-Aceves in the way in which a router obtains its own source tree from the trees reported by its neighbors. An important feature of all prior routing tree-based protocols is the fact that they are all based on communicating routing trees in terms of the lengths of the links and the identifiers of the nodes that form part of the routing trees.

SUMMARY OF THE INVENTION

In one embodiment of the present scheme, one or more labeled routing trees (LRTs) are produced at a router of a computer network according to a shortest path determination made over a partial topology graph of the network, which graph is produced according to knowledge of adjacent links of the router and one or more LRTs of neighboring routers. The LRTs of the router may be updated in response to receipt of routing state update messages, and such messages may include local link identifiers assigned by a head of a link to which the identifiers pertain, and node parameters of a tail of the link to which the local link identifiers pertain. The routing state update messages may be transmitted within the network: (i) in response to a new destination node being detected by an existing node within the network, (ii) in response to a destination becoming unreachable by a collection of the existing nodes, (iii) in response to the change in the cost of a path to at least one destination exceeding a threshold delta, and/or (iv) in situations where a routing loop may be encountered among two or more of the nodes of the network (e.g., at times when a path implied in the LRT of the router leads to a loop).

In another embodiment, the present routing protocol allows for distributing local link identifiers among nodes of a computer network within routing state update messages. These local link identifiers are preferably assigned by a head of a link to

which the identifiers pertain. The routing state update messages may include state parameters for nodes at tails of links to which the identifiers pertain. The routing protocol may further provide for distributing labeled routing trees of the nodes of the computer network among the nodes, for example wherein each node of the computers network may distribute its labeled routing trees to its neighboring nodes.

In such a routing protocol each node of the computer network preferably maintains one labeled routing tree per type of service offered in the network. These labeled routing trees may be generated according to a partial topology graph derived from the node's adjacent links and labeled routing trees of neighboring nodes, for example by applying a path selection algorithm to the partial topology graph. The labeled routing trees may then be updated in response to receipt of the routing state update messages. Such updates may be made according to either an optimum routing approach or a least overhead routing approach, depending upon which approach is used within the computer network.

For the optimum routing approach, routing state update messages are preferably transmitted: (i) when a routing state update message which causes a link to be added to a subject node's labeled routing trees is received, (ii) when a more recent routing state update message regarding an existing link in the subject node's labeled routing trees is received, and/or (iii) when a destination becomes unreachable by the subject node. In the least overhead routing approach, routing state update messages may be transmitted: (i) when a subject node finds a new destination or any neighbors of the subject node report same, (ii) when a destination becomes unreachable by the subject node or any of its neighbors, (iii) when the subject node finds that the change in the cost of a path to at least one destination exceeds a threshold delta, (iv) when a path implied in any of the labeled routing trees of the subject node leads to a loop, and/or a new successor chosen to a given destination has an address larger than an address of the subject node and a reported distance from the new successor to a particular destination is larger than a reported distance from a previous successor to that particular destination. Another embodiment provides a routing update message that includes information regarding

performance characteristics and addressing information for a link of a computer network and a node of the network at a tail end of the link. In some cases, the routing update may further include a time stamp assigned by a node at a head end of the link, a type of service vector, and addressing information for the node at a head end of the link. The performance characteristics preferably include link state parameters (e.g., link delay, link cost, available bandwidth, and/or reliability) of the link, while the type of service vector preferably specifies a type of service routing tree in which the link is being used by a node transmitting the routing update message. The addressing information of the link is preferably specified in the form of a local link identifier assigned to the link by the node at a head end of the link.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements and in which:

Figure 1 illustrates an example of an ad hoc wireless network with routers configured in accordance with the present routing protocol;

Figures 2a - 2d illustrate an example of a six-node wireless network with routers configured in accordance with the present routing protocol and the labeled routing trees developed at various ones of the routers;

Figures 3a - 3f illustrate an example of a six-node wireless network with routers configured in accordance with the present routing protocol and the routing state updates generated after certain link failures within the network; and

Figures 4a - 4f illustrate an example of how a link failure within a wireless network made up of a number of routers configured in accordance with the present routing protocol may not lead to the generation of new update messages by such routers when all such routers still have a path to all available destinations in the network.

DETAILED DESCRIPTION

A scheme for enabling routing of data packets in a computer network along preferred paths either on a hop-by-hop basis, or by specifying a source route efficiently by means of local identifiers is disclosed herein. Although discussed with reference to certain illustrated embodiments, upon review of this specification, those of ordinary skill in the art will recognize that the present scheme may find application in a variety of systems. Therefore, in the following description the illustrated embodiments should be regarded as exemplary only and should not be deemed to be limiting in scope.

In the present scheme, each packet being routed carries a routing operation code (ROC) instructing the receiving router which routing method to apply to forward the packet. A packet can be forwarded in any of the following forwarding modes: (a) a conventional hop-by-hop routing mode, which uses the routing table of the router; or (b) a source-routing mode, in which the entire source route is specified in the packet using local link identifiers instead of the addresses of relay routers as is common in prior source routing approaches. These forwarding modes are enabled using the present routing protocol, termed the Adaptive Internet Routing (AIR) protocol, which allows for the dissemination of link-state information and node-state information in the form of labeled routing trees (LRTs).

With AIR, a router sends updates to its neighbors regarding the links and nodes in its preferred paths to destinations. The links and nodes along the preferred paths from a source to each desired destination constitute an LRT that implicitly specifies the complete paths from the source to each destination. Each link is labeled with a local link identifier (LLID).

Each link in an LRT is directed and has a head-of-link node and a tail-of-link node. The head of the link labels the link with an LLID that differentiates that link from all other links from the same node to other neighbors. The LLID of a link is much smaller (in terms of the number of bits required to specify the LLID) than the address of a node.

Each router maintains an LRT for each type of service defined in the network (e.g., minimum-hop, minimum delay, and maximum bandwidth paths of the smallest number of hops). Each router also maintains a routing or topology graph that includes the state information about adjacent links and the LRTs reported by its neighbors. Each router computes its LRT for a given type of service based on the link and node information available in its topology graph.

A router reports changes to any of its LRTs to all its neighbors incrementally or atomically. The rules used to decide when a router should communicate changes to the state of a node or link can be based on optimum routing or least overhead routing approaches. The aggregation of adjacent links and routing trees reported by neighbors then constitutes the partial topology known by a router. To minimize the number of times the state of a link is communicated to neighbors, a router may assign a type vector to each link, with each bit of the vector indicating in which LRT the link is being used by the router.

AIR does not require backbones, the dissemination of complete cluster topology within a cluster, or the dissemination of the complete inter-cluster connectivity among clusters. Furthermore, AIR can be used with distributed hierarchical routing schemes proposed in the past for either distance-vector or link-state routing. See, e.g., L. Kleinrock and F. Kamoun, "Hierarchical Routing for Large Networks: Performance Evaluation and Optimization," *Computer Networks*, Vol. 1, pp. 155-174, 1977; M. Steenstrup (ed.), *Routing in Communication Networks*, Prentice-Hall, 1995; S. Murthy and J.J. Garcia-Luna-Aceves, "Loop-Free Internet Routing Using Hierarchical Routing Trees," *Proc. IEEE INFOCOM 97*, Kobe, Japan, April 7-11, 1997; and J. Behrens and J. J. Garcia-Luna-Aceves, "Hierarchical Routing Using Link Vectors," *Proc. IEEE INFOCOM 98*, San Francisco, California, March 29-April 2, 1998.

A router chooses its preferred paths for a given type of service using a local path-selection algorithm. One preferred path-selection algorithm for use according to the present scheme is a modification of the shortest-path first (SPF) algorithm. The result of running the path-selection algorithm over the topology graph is an LRT for a

given type of service that specifies, for each node in the LRT: the address of the head of the link incident to the node, the state parameters of the link, the state parameters of the node, and the LLID assigned to the link by the head of the link. Of course, other path selection algorithms may be used to provide similar outputs.

From its LRT, a router can compute a source route to a given destination. Because links are labeled with LLIDs assigned by the head of the links, a router can uniquely specify a source route to a destination using its own address followed by a sequence of LLIDs corresponding to its preferred path to the destination, rather than as a sequence of much larger network or link-level addresses. Although source routing and the use of local link identifiers have been used independently of one another in routing and bridging protocols in the past, AIR is the first routing protocol that disseminates the LLIDs of links, thus allowing a compact specification of source routes, making much more efficient use of the available bandwidth.

The present routing scheme is thus well suited for an ad hoc network that provides a seamless extension of the Internet Protocol (IP) to the ad hoc wireless environment. AIR will be described in terms of its operation in internet radios or IRs, which are wireless routers. However, it will be evident to those of ordinary skill in the art that AIR applies to computer networks and inter-networks that need not be based on wireless links for router interconnection. Figure 1 illustrates aspects of an exemplary ad hoc internet that will assist in understanding the remaining discussion.

Ad hoc network 10 may be considered as a number of sub-networks 12a, 12b, 12c, which provide an extension of the Internet 14 through a number of Internet Radios or IRs 16a-16i. Each IR 16a-16i is a wireless router with an assigned IP address and a medium access control (MAC) address. In general, IRs 16a-16i operate over one or a multiplicity of radio channels using spread-spectrum wireless communication techniques common in the art. For example, the IRs 16a-16i may operate in one of the unregulated UHF frequency bands, thereby obviating the need for operating licenses. Coupling of ad hoc network 10 to the Internet 14 is achieved through a router 18, which may be operated by an Internet Service Provider (ISP). As shown, a single ISP may

operate a LAN 20 to which multiple IRs are connected. In such a scheme, IRs 16a and 16b may act as "AirHeads", providing gateway service to Internet 14 via router 18. Some IRs, e.g., IRs 16d and 16e of Figure 1, may be associated with hosts, 22a, 22b and 22c, which can be accessed by any Internet user through ad hoc network 10. Like any router, each IR 16a-16i processes all messages, changes in the cost of an adjacent link, adjacent-link failures, and new-neighbor notifications one at a time and in the order in which it detects them.

Any IR 16a-16i in Figure 1 can consider another IR to be adjacent (we call such an IR a "neighbor") if there is radio connectivity between the two IRs and one IR, e.g., IR 16g, can receive and acknowledge packets from the other IR, e.g., IR 16h. Accordingly, a physical broadcast link connecting multiple IRs is mapped into multiple point-to-point bidirectional links defined for the same IRs. Each pair of adjacent IRs defines two point-to-point bidirectional links between them, one in each direction. Each point-to-point bidirectional link has a head node of the link and a tail node of the link.

The present routing scheme can be brought to practice together with the methods described in co-pending Application No. 09/248,738, entitled "Adaptive Communication protocol for Wireless Networks," filed February 10, 1999, and assigned to the Assignee of the present invention, incorporated herein by reference, for the assignment of logical link identifiers by one node to each of its one-hop neighbors. However, it will be evident to those of ordinary skill in the art that the present scheme can make use of any link-level service that enables a router to use a local identifier for each of its one-hop neighbors. The description of the present routing scheme thus assumes the existence of local identifiers for the links between a router and its immediate neighbors.

An underlying protocol, the above-noted neighbor protocol, assures that each IR 16a-16i detects within a finite time the existence of a new neighbor IR and the loss of connectivity with a neighbor IR. The neighbor protocol assumed in the present scheme can be brought to practice using link-layer retransmission strategies common in the art.

I. AIR Operation

In AIR, each IR reports to its neighbors the characteristics of every link and IR it uses to reach a destination in the ad hoc network. The set of links and IRs used by an IR in its preferred path to destinations is called the LRT of the IR. If multiple types of service (TOS) are defined in the network, an IR maintains one LRT per TOS. An IR therefore knows its adjacent links and the LRTs reported by its neighbors for each TOS, and the aggregation of an IR's adjacent links and the LRTs reported by its neighbors constitutes a partial topology graph. The links in the LRT and topology graph should be adjacent links or links reported by at least one neighbor. An IR may use the topology graph to generate its own LRT for each TOS. Each IR derives its LRT and routing table specifying the successor to each destination for each TOS by running a local path-selection algorithm for each TOS on its topology graph.

An IR communicates the updates it makes to its routing tree for any TOS. Because each IR communicates its routing tree for each TOS to its neighbors, the deletion of a link no longer used to reach a destination for a given TOS is implicit with the addition of the new link used to reach the destination and need not be sent explicitly as an update; an IR makes explicit reference to a failed link only when the deletion of a link causes the IR to have no paths to one or more destinations, in which case the IR cannot provide new links to make the deletion of the failed link implicit.

The basic update unit used in AIR to communicate changes to routing trees is the routing-state update (RSU), which describes the state of a link and the node at the end of the link. The head node of a link is the only IR that can report changes in the parameters of that link. RSUs are validated using sequence numbers, and each IR erases a link from its topology graph if the link is not present in the routing trees of any of its neighbors.

The operation of AIR can be described formally by pseudocode. As with any such description, however, there are various equivalent formulations that, upon review of the code provided, will be evident as being equivalent thereto to someone of ordinary

skill in the art. The pseudocode description is set forth at the end of this detailed description, and the following variables and expressions are used therein:

T : a constant defining the number of Types of Service (TOS) in the network

$\Delta[T]$: $\Delta[0], \Delta[1], \dots, \Delta[T-1]$ is a vector corresponding to thresholds to changes in the cost of a path to a destination for a given TOS.

AGEOUT_INTERVAL: a constant defining the time a failed link stays in the Topology Graph before being aged-out

ORA: a constant, TRUE if AIR is running under the Optimum Routing Approach

LORA: a constant, TRUE if AIR is running under the Least Overhead Routing Approach

TG_i : Topology Graph at router i

LRT_i : Current Labeled Routing Tree at router i

LRT_i' : Labeled Routing Tree created the last time an update message was generated by i

N_i : set of neighbors of router i

NS_i : set to TRUE if i has not sent its Labeled Routing Tree to a neighbor

M_i : set to TRUE if i must report changes to the Labeled Routing Tree

T_i : system time used to generate time stamps to RSU_i

$(u, v, llid, t, l[T], \{n\}, tos[T], del)$: An entry for link (u, v) in TG_i, LRT_i, TG_k^i , and LRT_k^i , where $k \in N_i$ and $\{l\} = \{llid, l[T], tos[T], del\}$.

u : Network address of the head of the link

v : Network address of the tail of the link

$llid$: Local link identifier

t : timestamp

$l[T]$: $l[0], l[1], \dots, l[T-1]$ is a vector corresponding to the performance parameters of the link. The parameter $l[0]$ corresponds to the cost of the link, the remaining parameters may be delay, bandwidth, and reliability of the link, etc.

$\{n\}$: The state parameters of router v

14

set $\text{tos}[T]: \text{tos}[0], \text{tos}[1], \dots, \text{tos}[T-1]$ is the Type of Service bit-vector. A bit x is to 1 when the link is added to the Labeled Routing Tree with TOS x .

del : Set to TRUE if the link cannot be used in the computation of the Labeled Routing Trees

$(d[T], \text{pred}[T], \text{suc}[T], d'[T], d''[T], \text{suc}'[T], \text{nbr})$: Variables assigned to a vertex v in TG_i ,

$\text{LRT}_i, \text{TG}_k^i$, and LRT_k^i , where $k \in N_i$

$d[T]$: Cost of the path $i \rightarrow v$ in Labeled Routing Tree x ; $\forall x \in [0; T-1]$

$\text{pred}[T]$: Predecessor (link) of vertex v in Labeled Routing Tree x ; $\forall x \in [0; T-1]$

1]

$\text{suc}[T]$: Next-hop towards vertex v in Labeled Routing Tree x ; $\forall x \in [0; T-1]$

$d'[T]$: Previous distance to v reported by $\text{suc}'[T]$

$d''[T]$: Cost of the path $i \rightarrow v$ the last time the cost of the path changed by $\Delta[T]$

$\text{suc}'[T]$: Previous successor towards v in Labeled Routing Tree x ; $\forall x \in [0; T-0]$

1]

nbr : Set to i if the cost of the path to v has increased but no update message must be generated

$(u, v, \text{llid}, t, l[T], \{n\}, \text{tos}[T])$: Entry in an update message (RSU)

u : Network address of the head of the link

v : Network address of the tail of the link

llid : Local link identifier

t : Time stamp assigned to RSU

$l[T]$: Vector corresponding to the performance parameters of the link

$\{n\}$: The state parameters of router v

$\text{tos}[T]$: Type of Service bit-vector

The pseudocode specifies the main procedures of AIR used to update the routing table and the link-state database at a router i for both an Optimum Routing Approach

(ORA) and a Least Overhead Routing Approach (LORA). Procedure NodeUp is executed when a router i starts up. The neighbor set of the router is empty initially.

If the neighbor protocol reports a new link to a neighbor k (procedure $\backslash\text{em NeighborUp}$), the router then runs Update with the appropriate message as input; the RSU in the message gets a current time stamp. The same approach is used for link failures (NeighborDown) and changes in the parameters of a link (LinkChange). When a router establishes connectivity to a new neighbor, the router sends its complete labeled routing tree to the neighbor (much like a distance vector protocol sends its complete routing table). The RSUs that are to be broadcast to all neighbors are inserted into MSG_i .

The procedure Update is executed when router i receives an update message from neighbor k or when the parameters of an outgoing link have changed. First, the topology graphs TG_i and TG_k^i are updated, then the labeled routing trees LRT_k^i and LRT_i are updated, which may cause the router to update its routing table and to send its own update message.

The state of a link in the topology graph TG_i is updated with the new parameters for the link if the routing-state update in the received message is valid, i.e., if the RSU has a larger time stamp than the time stamp of the link stored in TG_i .

The parameters of a link in TG_k^i are always updated when processing an RSU sent by a neighbor k , even if the link-state information is outdated, because they report changes to the labeled routing tree of the neighbor. A node in a labeled routing tree LRT_k^i for a given TOS can have only one link incident to it. Hence, when i receives an RSU for link (u,v) from k the current incident link (u', v) to v , $u \neq u'$, is deleted from TG_k^i .

The information of an RSU reporting the failure of a link is discarded if the link is not in the topology graph of the router.

A shortest-path algorithm (SPF) based on Dijkstra's SPF (procedure BuildShortestPathTree) is run on the updated topology graph TG_k^i to construct a new

labeled routing tree LRT_k^i , and then run on the topology graph TG_i to construct a new labeled routing tree LRT_i .

The incident link to a node v in router's i new labeled routing tree is different from the link in the current labeled routing tree LRT_i only if the cost of the path to v has decreased or if the incident link in LRT_i was deleted from the labeled routing trees of all neighbors.

A new labeled routing tree newLRT for a neighbor k , including the router's new labeled routing tree, is then compared to the current labeled routing tree LRT_k^i (procedure UpdateNeighborTree), and the links that are in LRT_k^i but not in newLRT are deleted from TG_k^i . After deleting a link (u, v) from TG_k^i the router sets $TG_i(u, v).del$ to TRUE if the link is not present in the topology graphs TG_x^i , $\forall x \in N_i$.

If a destination v becomes unreachable, i.e., there is no path to v in the new labeled routing tree newLRT, then RSUs will be broadcast to the neighbors for each link in the topology graph TG_i that have v as the tail node of the link and a link cost infinity.

This specification assumes that the Link Layer provides reliable broadcast of network-level packets and consequently update messages specify only incremental changes to the router's labeled routing tree instead of the complete labeled routing tree.

The new router's labeled routing tree newLRT is compared to the last reported labeled routing tree ($\{LRT_i\}$ for LORA and LRT_i for ORA) (procedure ReportChanges), and an update message that will be broadcast to the neighbors is constructed from the differences of the two trees. An RSU is generated if the link is in the new labeled routing tree but not in the current labeled routing tree, or if the parameters of the link have changed. For the case of a router running LORA, the labeled routing trees are only compared with each other if at least one of the four rules described in section V, below, is met, i.e., $M_i = \text{TRUE}$.

If the new router's labeled routing tree was compared against the last reported labeled routing tree then the router removes from the topology graph all the links that

are no longer used by any neighbor in their labeled routing trees (failed links are only removed from the topology graph by aging).

Finally, the current shortest-path tree LRT_k^i is discarded and the new one becomes the current labeled routing tree. The router's labeled routing tree is then used to compute the new routing table, using for example a depth-first search in the shortest-path tree.

II. Information Exchanged in AIR

Prior routing protocols based on topology information or distance information are based on the parameters of links exclusively. In contrast, AIR uses an update unit that conveys information about the performance characteristics and addressing information for a link and the node at the end of the link. More specifically, an RSU includes the following elements:

- a) A time stamp that validates the RSU;
- b) A type-of-service vector;
- c) The network address of the head node of the link;
- d) The network address of the tail node of the link;
- e) The link state parameters of the link between the two IRs; and
- f) The node state parameters of the tail of the link.

An update message sent by an IR contains at least one RSU. The time stamp of the RSU is assigned by the head of the link and should not be altered by any other node relaying the RSU in an update message. The type-of-service (TOS) vector is a bit vector specifying the TOS routing tree in which the link is being used by the node sending the RSU. The state parameters of a link are specified as a list of tuples, with each tuple consisting of a type and a content. There are two classes of state parameters for a link: performance parameters and addressing parameters. The performance of a link can be characterized in terms of its delay, cost, bandwidth, and reliability, for example. An addressing parameter specifies an identifier assigned to the link. An example of such an identifier in the present scheme is the local link identifier (LLID) assigned to the link by

the head of the link. The state parameters of the tail of a link may include, for example, the remaining battery life of the node.

III. Information Stored in AIR

Figures 2 a - 2d illustrate the fact that IRs running AIR need maintain only partial topology information. These illustrations provide an example of a six-node wireless network (each node labeled a - f, respectively). For simplicity, these Figures assume that a single link parameter is used to characterize a link in one of its directions, called the "cost" of the directed link. For example, link 32c may have a cost of 5 in the direction b to c, but only a cost of 1 in the direction c to b. In other examples, nodes may be coupled by multiple links between them, each having the same or different costs. Figures 2b through 2d show the selected topology according to AIR at the IRs indicated with filled circles. Solid lines represent the links that are part of the labeled routing tree of the respective IR. Arrowheads on links indicate the direction of the link stored in the IR's topology graph. IR a's labeled routing tree shown in Figure 2b is formed by the labeled routing trees reported by its neighbors b and c, and the links for which IR a is the head node (namely links (a, b) and (a, c)). Similarly, Figure 2c shows the LRT for IR b and Figure 2d that for IR c. From the figures, the savings in storage requirements should be clear, even for this few-node network.

The information maintained by an IR to participate in AIR includes a topology graph, an LRT for each TOS defined in the network, a routing table, and an adjacent-link table. The record entry for the link from u to v in the topology graph consists of the tuple (u, v, t, {l}, {n}), where u and v are the network addresses of the head and tail of the link, respectively, t is the most recent time stamp received for link (u, v), {l} is a sequence of type-value pairs specifying link parameters, and {n} is a sequence of type-value pairs specifying node parameters. A link parameter used in the present scheme is the LLID of the link.

The routing table specifies, for each destination and for each TOS, the next IR in the path to the destination and the distance to that destination based on the distance metric used for the TOS.

The cost of a failed link is considered to be infinity for any TOS. There are various ways in which costs may be assigned to links for a given TOS known in the art. For example, the cost of a link could simply be the number of hops, or the addition of the latency over the link plus some constant bias.

IV. Validating Updates

Because of delays in the IRs and links of an internetwork, update messages sent by a IR may propagate at different speeds along different paths. Therefore, a given IR may receive an RSU from a neighbor with stale link-state information, and a distributed termination-detection mechanism is necessary for a IR to ascertain when a given RSU is valid and avoid the possibility of RSUs circulating forever. AIR uses time stamp to validate RSUs. An IR either maintains a clock that does not reset when the IR stops operating, or asks its neighbors for the oldest known time stamp after it initializes or reboots.

An IR receiving an RSU accepts the RSU as valid if the received RSU has a larger time stamp than the time stamp of the RSU stored from the same source, or if there is no entry for the link in the topology graph and the RSU is not reporting an infinite cost. Link-state information for failed links are the only RSUs erased from the topology graph due to aging (which may be on the order of an hour or so (or other time period) after having processed the RSU). RSUs for operational links are erased from the topology graph when the links are erased from the routing trees of all the neighbors.

It is noted that, because RSUs for operational links never age out, there is no need for the head node of a link to send periodic RSUs to update the time stamp of the link. This is important, because it means that AIR does not need periodic update messages to validate link-state information like OSPF, J. Moy, "OSPF Version 2," RFC

1583, Network Working Group, March 1994, and all prior protocols based on sequence numbers or time stamps.

V. Exchanging Update Messages

An IR sends RSUs in two different ways: (a) Following an optimum routing approach; and (b) following a least overhead approach. The optimum routing approach is well suited for networks with fairly static topologies. The least overhead approach is tailored for networks with dynamic topologies due to IR mobility. Which approach should be executed in IRs can be defined by an IR configuration parameter.

As indicated in the pseudocode description, according to the optimum routing approach an IR sends RSUs about a link in the following cases: (a) when an RSU is received for the link causing the link to be added to the IR's routing tree, (b) when the link is already in the IR's routing tree and a more recent RSU is received for the link, (c) when an RSU reporting the failure of the link results in no path to the tail of the link, and (d) when a failed link is not in the IR's routing tree and there is no path to the tail of the link.

In contrast, according to the least overhead approach an IR sends RSUs according to the following rules:

- (a) The IR finds a new destination, or any of its neighbors reports a new destination.
- (b) The IR finds that the change in the cost of a path to at least one destination exceeds a threshold delta, or at least one destination becomes unreachable to the IR or any of its neighbors.
- (c) A path implied in the LRT of the IR leads to a loop.
- (d) The IR sends an RSU when: (i) The new successor chosen to a given destination has an address larger than the address of the IR; and (ii) the reported distance from the new chosen successor n to a destination j is longer than the reported distance from the previous

successor to the same destination. However, if the link from the IR to j fails and n is a neighbor of j , no update message is needed regarding j or any destination whose path from the IR involves j .

Each time an IR processes an update message from a neighbor, it updates that neighbor's LRT and traverses that tree to determine for which destinations its neighbor uses the IR as a relay in its preferred paths. The IR then determines if it is using the same neighbor as a relay for any of the same destinations. A routing loop is detected if the IR and neighbor use each other as relay to any destination, in which case the loop must be broken and the IR must send an update message with the corresponding changes.

We observe that, in any routing loop among IRs with unique addresses, one of the IRs must have the smallest address in the loop; therefore, if an IR is forced to send an update message when it chooses a successor whose address is larger than its own, then it is not possible for all IRs in a routing loop to remain quiet after choosing one another, because at least one of them is forced to send an update message, which causes the loop to break when IRs update their LRTs.

To ensure that AIR works correctly with least overhead routing and incremental updates specifying only changes to an LRT, an IR must remember the LRT that was last notified to its neighbors. If any of the rules for update notification in least overhead routing are satisfied, the IR must do one of two things: (a) If the LRT includes new neighbors than those present in the LRT that was last updated, then the IR must send its entire LRT in its update, so that new neighbors learn about all the destinations the IR knows; (b) if the two LRTs imply the same neighbors, the IR sends only the updates needed to obtain the new LRT from the old one.

To ensure that AIR stops sending update messages, a simple rule can be used to determine which IR must stop using its neighbor as a relay, such a rule can be, for example, "the IR with the smaller address must change its path."

The rules for update-message exchange according to least overhead routing stated above assume that an update message is sent reliably to all the neighbors of an IR.

The following example illustrates a scenario in which the last rule is needed to prevent permanent loops. Consider the six-node wireless network 30 shown in Figure 2a. In this example, IRs are given identifiers that are lexicographically ordered, i.e., "a" is the smallest identifier and "f" is the largest identifier in the graph. All links and nodes are assumed to have the same propagation delays, all the links but links (a, b) and (b, c) have unit cost, and $\delta = \infty$. Figures 2b through 2d show the LRTs according to AIR at the IRs indicated with filled circles for the network topology depicted in Figure 2a. Arrowheads on solid lines indicate the direction of the links stored in the IR's LRT.

Figures 3a - 3f now depict the sequence of events triggered by the execution of AIR in the example network after the failures of links (c, d) 32g and (b, e) 32d. The figures show the RSUs (in parentheses) generated by the node with filled circle, which RSUs are transmitted in an update message to the node's neighbors. The third element in an RSU corresponds to the cost of the link (a RESET has cost infinity). As shown in Figure 3b, node c transmits an RSU after processing the failure of link (c, d) 32g; the distance from the new successor b to d and f is longer than from the previous successor d. When link (b, e) 32d fails (see Figure 3d), node b realizes that the destinations d, e, and f are unreachable and generates an RSU reporting the failure of the link connecting to the head of the subtree of the LRT that becomes unreachable. The RSU from b triggers the RSUs that allow nodes a, b, and c to realize that there are no paths to d, e, and f (Figures 3e and 3f). A similar sequence of events takes place at the other side of the network partition (not shown).

As another example of the operation of AIR, consider the seven-node wireless network 40 shown in Figure 4a. All links and nodes are assumed to have the same propagation delays, all the links have unit cost, $\delta = \infty$. Figures 4b through 4d show the LRTs produced according to AIR at the IRs indicated with filled circles for the network topology depicted in Figure 4a. Arrowheads on solid lines indicate the direction of the links stored in the IR's LRT. When the link (f, g) fails (Figure 4e), the neighbor protocol at node f triggers the execution of procedure NeighborDown, the link (d, g) is inserted into f's LRT but no update message is generated because f's new successor

towards g has an address smaller than f and destination g is a neighbor of the new successor. Figure 4f shows the new LRT of node d after the failure of link (d, g).

Because d has an address smaller than the new successor towards g it is required to send an update message reporting the new link added to the LRT. Nodes c, e, and f do not generate any update message after processing d's message because there exist a path to all destinations in the network and no routing loop was formed. This example thus illustrates how link failures may not cause the generation of update messages by nodes that have the failed link in their LRTs as long as the nodes have a path to all destinations.

VI. Obtaining LRTs and Source Routes

Obtaining the LRTs and source routes in the present scheme is done with a very simple modification to Dijkstra's SPF algorithm that is run on the topology graph of an IR. The modifications to SPF consist of checking for the proper bit to be set in the TOS bit vector of a link so that only those links being used for the required TOS are considered in preferred paths, and accumulating the source route in terms of LLIDs as the topology graph is traversed.

Given the topology graph, the algorithm proceeds as follows:

1. Initialize:

Value of TOS bit vector = T

Set IR-set = {root}, where root is the IR running the algorithm.

Dist-to-root = 0

Route-to-root = root

for all IRs other than root set

Dist-to-IR = infinity

Route-to-IR = root

for all IRs neighbors of root

If TOS bit vector = T then

Dist-to-IR = cost of link to IR

Route-to-IR = Route-to-IR \cup LLID of link to IR

2. Find next IR in IR-set:

Find an IR x not in set such that:

Dist-to-IR = Minimum{Dist-to-IR not in IR-set

considering only links with TOS bit vector = T}

Augment IR-set = IR-set \cup x

Stop if IR-set contains all IRs

3. Change chosen distance and path:

For each IR y not in IR-set do:

Dist-to-y = Minimum{ Dist-to-y, Dist-to-z + cost of link(z,y)

with z in IR-set and TOS bit vector of (z,y) = T}

Route-to-y = Route-to-z \cup LLID of (z,y)

4. Repeat Step 2.

Thus a scheme for enabling routing of data packets in a computer network has been described. Although the foregoing description and accompanying figures discuss and illustrate specific embodiments, it should be appreciated that the present invention is to be measured only in terms of the claims that follow the example of a pseudocode listing for the routing protocol set forth below:

CLAIMS

What is claimed is:

1. A method, comprising producing one or more labeled routing trees (LRTs) at a router of a computer network according to a shortest path determination made over a partial topology graph of the network, which graph is produced according to knowledge of adjacent links of the router and one or more LRTs of neighboring routers.
2. The method of claim 1 wherein the LRTs of the router are updated in response to receipt of routing state update messages.
3. The method of claim 2 wherein the routing state update messages include local link identifiers assigned by a head of a link to which the identifiers pertain.
4. The method of claim 3 wherein the routing state update messages further include node parameters of a tail of the link to which the local link identifiers pertain.
5. The method of claim 3 wherein the routing state update messages are transmitted within the network: (i) in response to a new destination node being detected by an existing node within the network, (ii) in response to a destination becoming unreachable by a collection of the existing nodes, and/or (iii) in response to the change in the cost of a path to at least one destination exceeding a threshold delta, (iv) in situations where a routing loop may be encountered among two or more of the nodes of the network.
6. The method of claim 5 wherein situations where a routing loop may be encountered include times when a path implied in the LRT of the router leads to a loop.
7. A routing protocol, comprising distributing local link identifiers among nodes of a computer network within routing state update messages.
8. The routing protocol of claim 7 wherein the local link identifiers are each assigned by a head of a link to which the identifiers pertain.
9. The routing protocol of claim 7 wherein the routing state update messages include state parameters for nodes at tails of links to which the identifiers pertain.

10. The routing protocol of claim 7, further comprising distributing labeled routing trees of the nodes of the computer network among the nodes.

11. The routing protocol of claim 10 wherein each node of the computers network distributes its labeled routing trees to its neighboring nodes.

12. The routing protocol of claim 11 wherein each node of the computer network maintains one labeled routing tree per type of service offered in the network.

13. The routing protocol of claim 11 wherein each node generates its labeled routing trees according to a partial topology graph derived from the node's adjacent links and labeled routing trees of neighboring nodes.

14. The routing protocol of claim 13 wherein the labeled routing trees of a node are generated by applying a path selection algorithm to the partial topology graph.

15. The routing protocol of claim 14 wherein each node of the computer network maintains one labeled routing tree per type of service offered in the network.

16. The routing protocol of claim 15 wherein labeled routing trees of nodes in the computer network are updated in response to receipt of the routing state update messages.

17. The routing protocol of claim 16 wherein labeled routing trees are updated according to whether an optimum routing approach or a least overhead routing approach is used within the computer network.

18. The routing protocol of claim 17 wherein in the optimum routing approach, routing state update messages are transmitted: (i) when a routing state update message which causes a link to be added to a subject node's labeled routing trees is received, (ii) when a more recent routing state update message regarding an existing link in the subject node's labeled routing trees is received, and/or (iii) when a destination becomes unreachable by the subject node.

19. The routing protocol of claim 17 wherein in the least overhead routing approach, routing state update messages are transmitted: (i) when a subject node finds a new

destination or any neighbors of the subject node report same, (ii) when a destination becomes unreachable by the subject node or any of its neighbors, (iii) when the subject node finds that the change in the cost of path to at least one destination exceeds a threshold delta, (iv) when a path implied in any of the labeled routing trees of the subject node leads to a loop, and/or a new successor chosen to a given destination has an address larger than an address of the subject node and a reported distance from the new successor to a particular destination is larger than a reported distance from a previous successor to that particular destination.

20. A routing update message, comprising information regarding performance characteristics and addressing information for a link of a computer network and a node of the network at a tail end of the link.

21. The routing update message of claim 20 wherein the routing update message further comprises a time stamp assigned by a node at a head end of the link.

22. The routing update message of claim 21 wherein the performance characteristics include link state parameters of the link.

23. The routing update message of claim 21 wherein the routing update message further comprises a type of service vector.

24. The routing update message of claim 21 wherein the routing update message further comprises addressing information for a node at a head end of the link.

25. The routing update message of claim 23 wherein the type of service vector specifies a type of service routing tree in which the link is being used by a node transmitting the routing update message.

26. The routing update message of claim 20 wherein the performance characteristics of the link comprise link delay, link cost, available bandwidth, and/or reliability.

27. The routing update message of claim 20 wherein the addressing information of the link comprises a local link identifier assigned to the link by a node at a head end of the link.

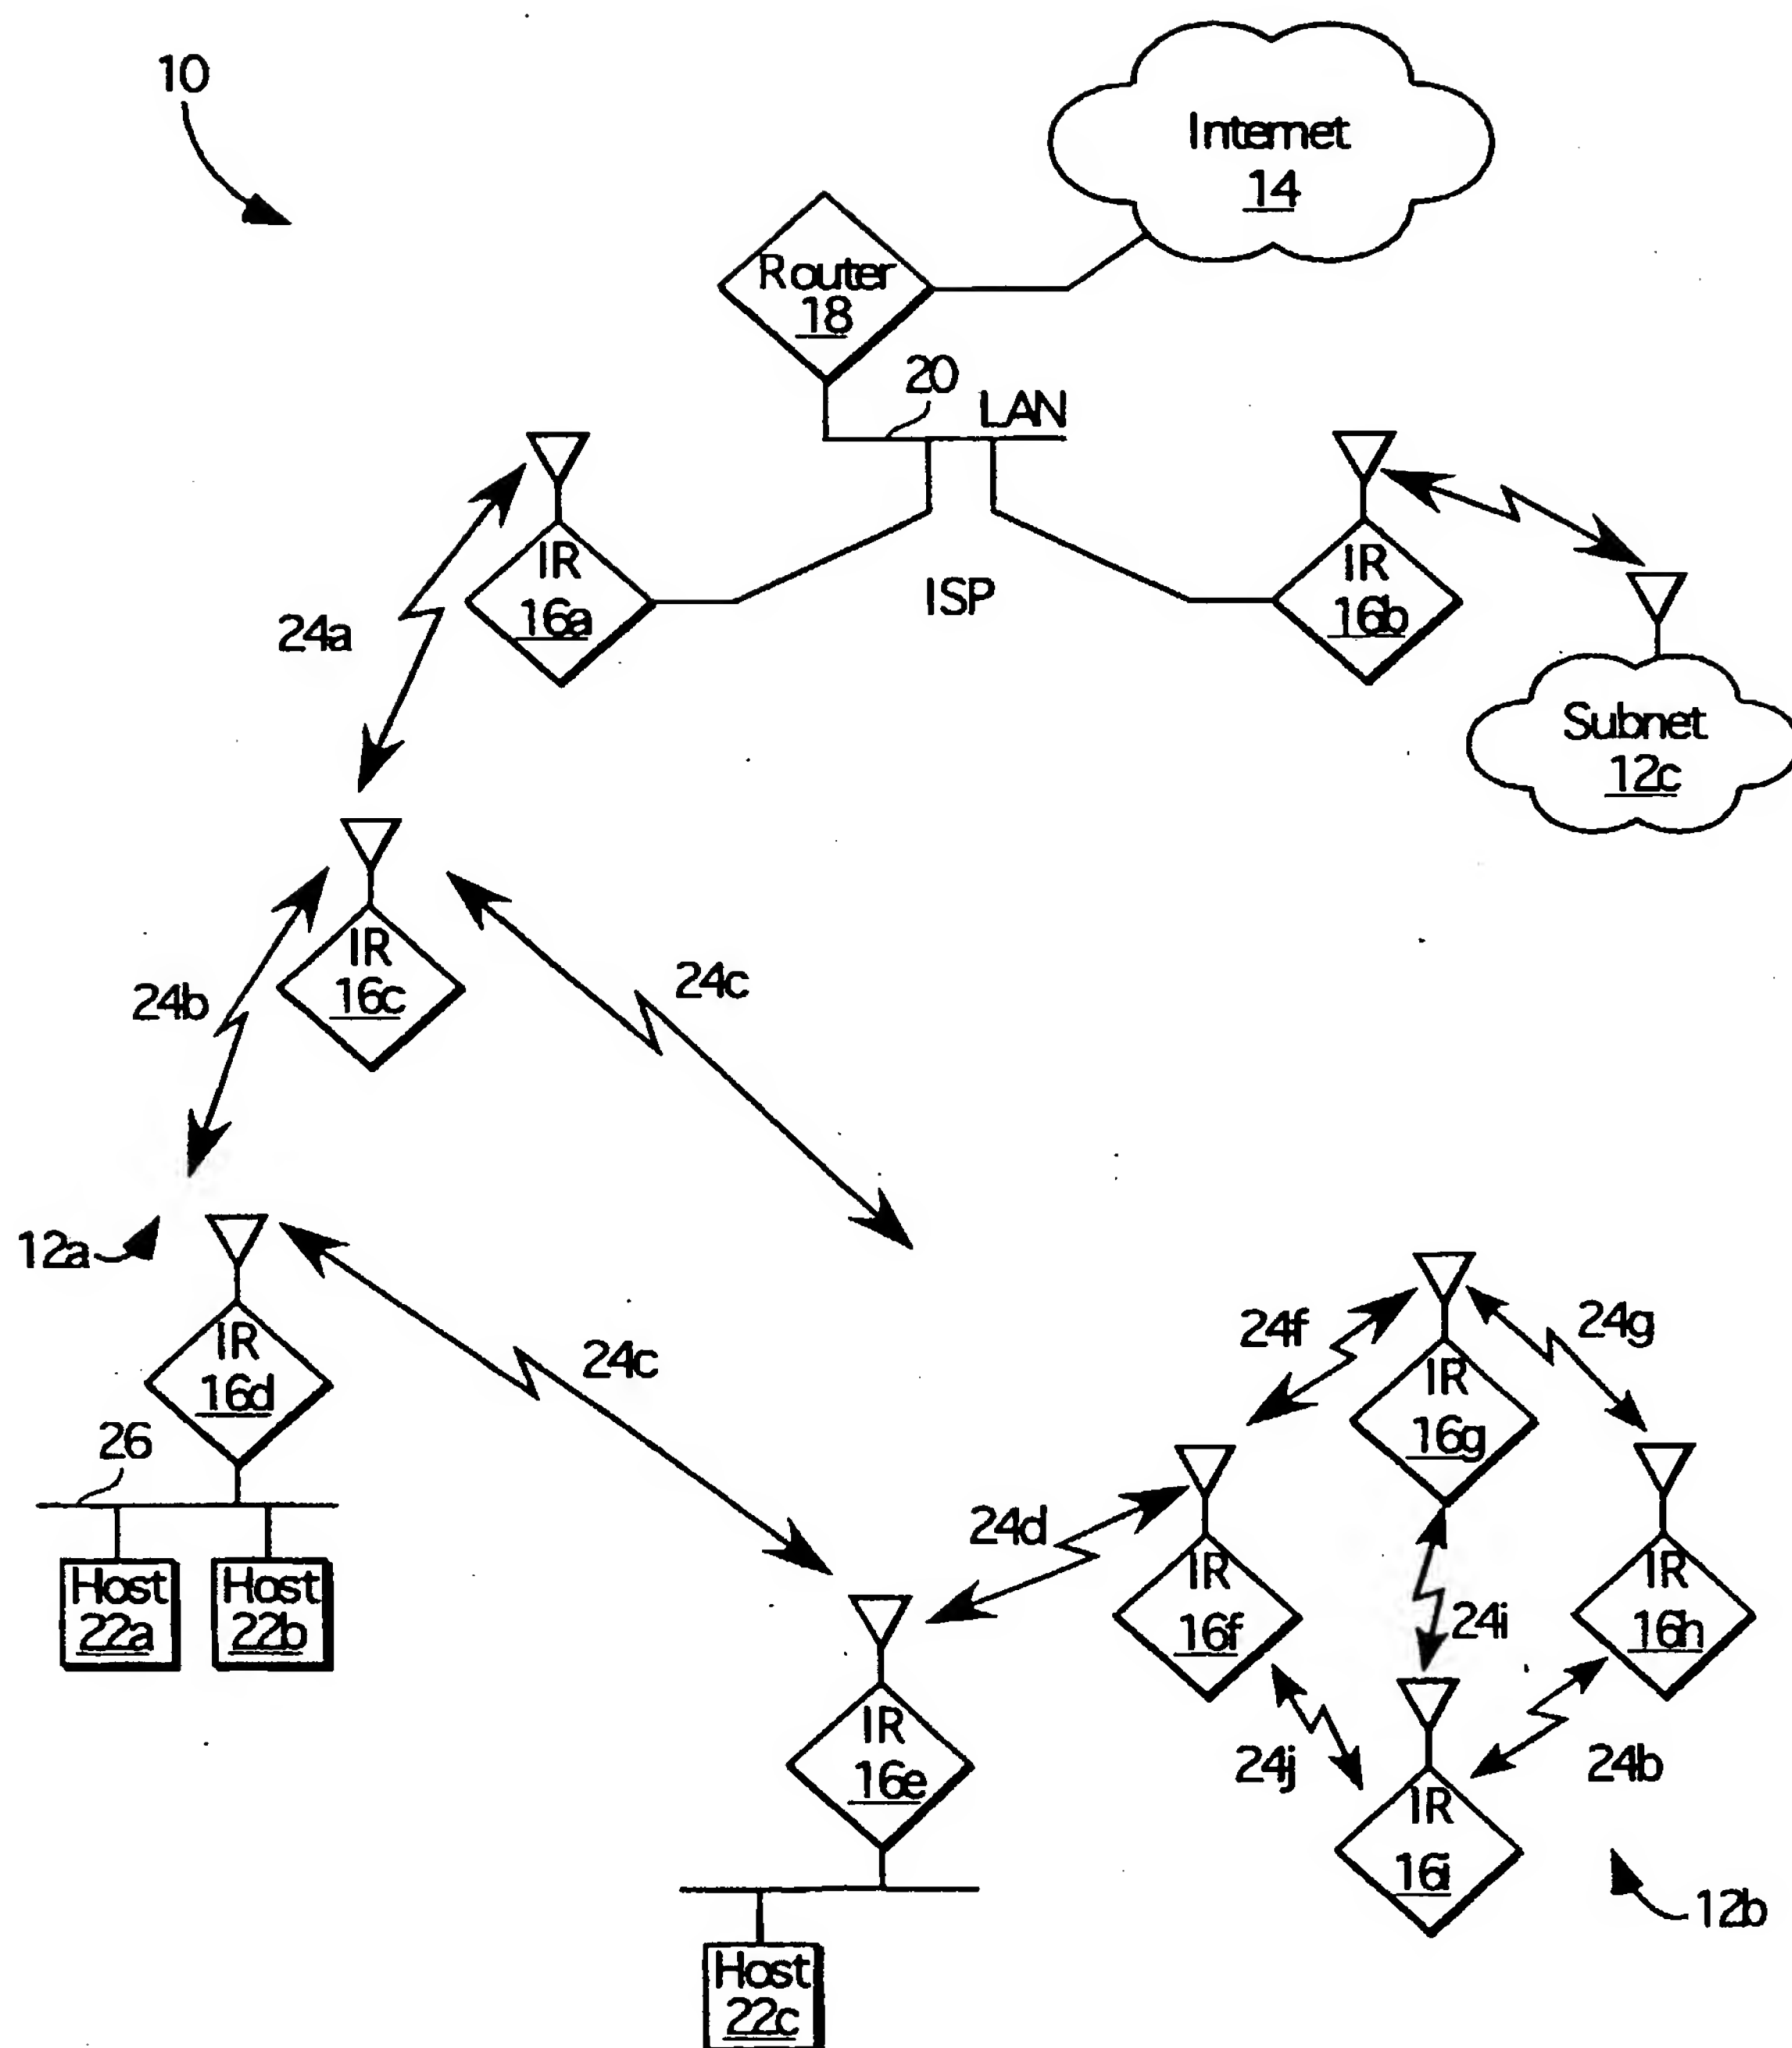


FIG. 1

2/4




FIG. 2A




FIG. 2B




FIG. 2C






FIG. 2D




(a)

FIG. 3A




(b)

FIG. 3B




(c)

FIG. 3C




(d)

FIG. 3D



(e)

FIG. 3E



(f)

FIG. 3F




FIG. 4A




FIG. 4B




FIG. 4C




FIG. 4D




FIG. 4E




FIG. 4F

INTERNATIONAL SEARCH REPORT

Int onal Application No

PCT/US 00/41180

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WCNC. 1999 IEEE WIRELESS COMMUNICATIONS AND NETWORKING CONFERENCE, vol. 03, 24 September 1999 (1999-09-24), pages 1308-1312, XP002168435 Piscataway, NJ, USA, IEEE, USA ISBN: 0-7803-5668-3 IEEE Catalog Number: 99TH8466 page 1308, right-hand column, line 10 -page 1309, left-hand column, line 38 -----	1-27



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the International filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

29 May 2001

Date of mailing of the international search report

15/06/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Siebel, C